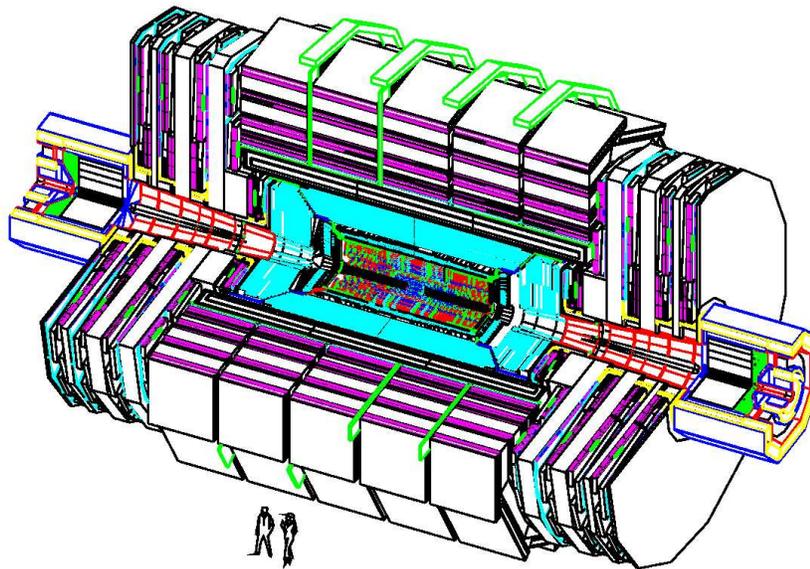


Heavy Ion Physics at the LHC with the Compact Muon Solenoid Detector

M. Ballintijn⁷⁾, O. Barannikova²⁾, R.R. Betts²⁾, B. E.Bonner⁹⁾, W. Busza⁷⁾, M. Calderón de la Barca Sánchez¹⁾, D. Cebra¹⁾, P. Constantin⁵⁾, G. Eppley⁹⁾, E. García²⁾, N. George³⁾, O. Grachov⁴⁾, S.V. Greene¹⁰⁾, D. Hofman²⁾, R. Hollis²⁾, A. Iordanova²⁾, J. Kapusta⁸⁾, G. J. Kunde⁵⁾, C. Loizides⁷⁾, W. Llope⁹⁾, C.F. Maguire¹⁰⁾, A.C. Mignerey⁶⁾, C. Mironov⁵⁾, M. Murray⁴⁾, G. Mutchler⁹⁾, E. Norbeck³⁾, Y. Onel³⁾, C. Roland⁷⁾, G. Roland⁷⁾, S.J. Sanders⁴⁾, G.S.F Stephans⁷⁾, M.B. Tonjes⁶⁾, J. Velkovska¹⁰⁾, R. Vogt¹⁾, B. Wyslouch⁷⁾, P. Yepes⁹⁾



¹⁾ University of California at Davis, Department of Physics, Davis, California 95616

²⁾ University of Illinois at Chicago, Department of Physics, Chicago, Illinois 60607

³⁾ University of Iowa, Physics and Astronomy Department, Iowa City, Iowa 52242

⁴⁾ University of Kansas, Department of Physics and Astronomy, Lawrence, Kansas 66045

⁵⁾ Los Alamos National Laboratory, Los Alamos, New Mexico, 87545

⁶⁾ University of Maryland, Chemistry Department, College Station, Maryland 20742

⁷⁾ Massachusetts Institute of Technology, Laboratory for Nuclear Science, Cambridge, Massachusetts 02139

⁸⁾ University of Minnesota, School of Physics and Astronomy, Minneapolis, Minnesota 55455

⁹⁾ Rice University, Bonner Nuclear Laboratory MS315, Houston, Texas 77005

¹⁰⁾ Vanderbilt University, Department of Physics and Astronomy, Nashville, Tennessee 37235

Contents

1	Overview	4
2	Heavy Ion Physics at the LHC	6
2.1	Science Questions	6
2.2	Learning about QCD from Heavy Ion Collisions at LHC	6
2.2.1	Initial Temperature and Quarkonia Yields	6
2.2.2	Transport Properties of the Medium and Jet Physics	7
2.2.3	Thermalization and Elliptic Flow	8
2.2.4	The Initial State and Global Observables	9
2.2.5	Event Characterization	9
2.2.6	Parton Distribution Functions	10
2.2.7	Forward Physics	10
2.2.8	Proton-Nucleus Collisions	11
2.2.9	Ultra-Peripheral Collisions	11
2.2.10	Proton-Proton Physics	11
2.2.11	Expected Rates for Hard Probes	12
2.3	Physics Interests of the US groups	13
2.4	Physics Schedule	13
3	The CMS Detector	16
3.1	Inner Si Tracker	17
3.2	Calorimeters	17
3.3	Muon System	18
3.4	Detectors in the Forward Region	18
4	CMS Detector Performance for Heavy Ion Collisions	19
4.1	Tracking Performance	19
4.2	Muon Reconstruction	22
4.3	Jet Reconstruction	23
4.4	Calorimetry	25
4.5	Comparison to ALICE	25
4.6	Summary	26
5	The CMS DAQ and Trigger System	27
5.1	Trigger Strategy for Heavy Ion Collisions	27
5.2	Trigger Channels in Heavy Ion Collisions	28

6	Overview of the Data Acquisition System	31
6.1	Data Acquisition System Architecture	31
6.2	Filter Unit	32
6.3	The HLT Application	32
6.4	DAQ Deployment Schedule	32
7	Level-1 Trigger	33
7.1	Introduction	33
7.2	Minimum Bias Triggers	33
7.3	Jets at Level-1	34
7.4	Muons at Level-1	35
7.5	Level-1 Outlook	35
8	High Level Trigger Performance	36
8.1	Event Size	36
8.2	Trigger Timing Studies	37
8.2.1	Jet Finder Timing	38
8.2.2	Muon Finder Timing	38
8.3	Trigger Signal Rates	38
8.4	Trigger Background Rates	38
8.5	Output Rates	39
8.6	Benefits of HLT running at design luminosity	40
8.7	Low Luminosity Scenario	41
8.8	Summary	41
9	Funding Request	43
10	Personnel, Schedule and Management Plan	46
11	Appendices	48
A	Profiles of the collaborating groups	48
A.1	University of California at Davis	48
A.2	University of Illinois at Chicago	48
A.3	University of Iowa	49
A.4	University of Kansas	49
A.5	Los Alamos National Laboratory	49
A.6	University of Maryland	49
A.7	Massachusetts Institute of Technology	50

A.8	University of Minnesota	50
A.9	Rice University	50
A.10	Vanderbilt University	50
A.11	Acknowledgment of non-US Participants	51
B	Operations Budget	53
B.1	Collaboration Costs	53
B.2	Additional Travel Costs for US Institutions	53
C	Computing for Heavy Ions	54
C.1	The CMS Computing Model	54
C.2	Offline Computing for Heavy Ion Runs	54
C.2.1	Heavy Ion Data Flow	55
C.2.2	Heavy Ion Data Storage Requirements	55
C.2.3	Heavy Ion Computing Requirements	55
C.3	Heavy Ion Computing Purchase Plan	55
D	Zero Degree Calorimeter	57
D.1	Physics Capabilities of the ZDC	58
D.2	ZDC Maintenance Costs	58
D.3	Status of the ZDC Project	59

1 Overview

Heavy ion collisions at the Large Hadron Collider (LHC) provide an opportunity for an unprecedented expansion of the study of Quantum Chromodynamics (QCD) in systems with extremely high energy density.

Data collected by the four experiments at the Relativistic Heavy Ion Collider provide insight into what can be expected at the LHC. The results suggest that in heavy ion collisions at $\sqrt{s_{NN}} = 200$ GeV an equilibrated, strongly-coupled partonic system is formed. There is strong evidence that this dense partonic medium is highly interactive, perhaps best described as a quark-gluon fluid, and is also almost opaque to fast partons. In addition, many surprisingly simple empirical relationships describing the global characteristics of particle production have been found. An extrapolation to LHC energies suggests that the heavy ion program has significant potential for major discoveries. Similar to the expectations for high energy physics, heavy ion studies at the LHC will either confirm and extend the theoretical picture emerging from lower beam energies or challenge and redirect our understanding of strongly interacting matter at extreme densities. This will be accomplished both by extending existing studies over a dramatic increase in energy and also by bringing to bear a broad range of novel probes (such as high p_T jets and photons, Z^0 bosons, the Υ states, D and B mesons, and high-mass dileptons) which are accessible only at LHC energies.

In this proposal, we will show that the Compact Muon Solenoid (CMS) detector provides unique capabilities for focused measurements that exploit the new opportunities at the LHC. These measurements will directly address the fundamental science questions in the field of high density QCD. This potential was recognized by the CMS collaboration which has included a heavy ion group since its inception. The US component of that group played a key role in expanding and developing the physics program and now leads the overall effort. The apparatus provides unprecedented coverage for tracking and both electromagnetic and hadronic calorimetry combined with high precision muon identification. The detector is read out by a very fast data acquisition system and allows for sophisticated triggering.

In Section 2 we will describe the science case for heavy ion physics at LHC. Based on the knowledge gained in the first five years of RHIC running, and our current theoretical understanding, we provide examples of measurements that will address the fundamental science questions in this field. A description of the CMS detector (Section 3) and studies of its performance in the high multiplicity heavy ion environment (Section 4) illustrate that it excels in exactly those categories which the experience at RHIC indicates will be most critical.

The key component in exploiting the CMS capabilities in heavy ion collisions is the trigger system, which is crucial in accessing the rare probes expected to yield the most direct insights into the properties of high density strongly interacting matter. The general trigger strategy and overall structure of the trigger hardware are introduced in Sections 5 and 6. More details of the Level-1 and Higher Level Triggers are given in Sections 7 and 8. Our studies have led to a unique trigger strategy for CMS Pb+Pb running, where the event selection will be performed in a very large online CPU farm running offline reconstruction algorithms. The implementation of this strategy forms the core of the proposed US contribution to CMS heavy-ion operations.

Finally, the requested funding and the proposed schedule and management plan are described in Sections 9 and 10. This proposal asks for a total of 2 M\$ for FY2007 through FY2010. The bulk of this funding will be used to purchase computer hardware for use in the CMS trigger, with small amounts needed for establishing a trigger development facility in the US. It should be noted that this contribution represents about 10% of the overall cost of the full DAQ system and 0.4% of the total cost of the CMS experiment.

With the exception of 0.25 FTE technician in FY2007 for the initial setup of the trigger facility, no personnel are supported under this proposal. Instead all of the US CMS heavy ion physicists (profiled in Appendix A) will be funded by their individual group budgets. Additional operating costs to cover collaboration fees as well as the anticipated impact of this project on the groups' travel expenses are discussed in Appendix B. The capital cost of off-line computing to support CMS heavy ion physics at US institutions is discussed in Appendix C. Additional hardware projects to construct the ZDC (see Appendix D) and CASTOR detectors are also part of the US CMS heavy ion effort but are funded through other sources.

In summary, The US CMS heavy ion group is actively engaged in preparing to exploit the unique physics potential using CMS to study heavy ion collisions at the LHC. Group members have extensive experience in the relevant topics from several of the detectors at RHIC as well as experiments at other facilities. Consultation and cooperation

with the strong CMS high energy groups present at many of the institutions participating in the heavy ion effort adds strength to both programs. With the exception of the Zero Degree Calorimeter, whose funding is not included in this proposal, the US heavy ion group within CMS is not directly involved in the construction of the apparatus. Our effort is focused on designing the analysis tools to make the most efficient use of the detector capabilities and, equally importantly, designing the trigger and data acquisition systems needed to optimize the physics output of the allocated beam time. With the broad suite of detector capabilities, and the expertise with the CMS heavy ion group, the experiment is well positioned to take advantage of the physics opportunities at LHC.

2 Heavy Ion Physics at the LHC

2.1 Science Questions

In this section we will describe the science motivation for studying heavy ion collisions at the LHC. The ultimate goal in the study of relativistic heavy ion collisions is to experimentally explore the structure of the vacuum under varying conditions. QCD predicts that at high temperatures, but zero density of net quarks, the vacuum will undergo a phase change to a state that is most easily described in terms of quark and gluon degrees of freedom, the Quark-Gluon Plasma (QGP). Lattice QCD calculations show a cross-over to this different state at a temperature of around 170 MeV, corresponding to energy densities of 0.6 GeV/fm³.

A list of key science questions for the experimental study of heavy ion collisions has been formulated in the NSAC Long Range Plan. Quoting from the plan, these questions were:

- In relativistic heavy-ion collisions, how do the created systems evolve? Does the matter approach thermal equilibrium? What are the initial temperatures achieved?
- Can signatures of the deconfinement phase transition be located as the hot matter produced in relativistic heavy ion collisions cools? What is the origin of confinement?
- What are the properties of the QCD vacuum and what are its connection to the masses of the hadrons? What is the origin of chiral symmetry breaking?
- What are the properties of matter at the highest energy densities? Is the basic idea that this is best described using fundamental quarks and gluons correct?

Since these questions were posed in 2002, further studies at RHIC have provided additional information on the topics of initial state effects on high p_T suppression, the applicability of hydrodynamics, and the energy loss and thermalization of heavy quarks. These new data provide additional guidance on the most promising experimental directions at the LHC for addressing the science questions mentioned above. In addition, there is increasing evidence that the dense QCD medium evolves out of a universal high-density QCD state, the Color Glass Condensate (CGC), a dense, but weakly coupled system of gluons. Further data are needed to confirm the role of the CGC in the evolution of high energy hadronic collisions and to extend our understanding of this state. In the following, we will let these questions guide our description of the science case for the LHC and we will, in particular, describe what unique capabilities CMS has to address these questions.

2.2 Learning about QCD from Heavy Ion Collisions at LHC

At the LHC, the energy densities of the thermalized matter are predicted to be 20 times higher than at RHIC, implying a doubling of the initial temperature [1]. The higher densities of the produced partons result in more rapid thermalization and, consequently, the time spent in the quark-gluon plasma phase increases by almost a factor of three [1]. These dramatically different conditions may allow the hot, dense system to reach the weakly interacting, ideal gas quark-gluon plasma, or some other different state, in contrast to the strongly interacting plasma believed to be created at RHIC [2]. Below we describe how various measurements with CMS can elucidate the properties of this hot and dense medium.

2.2.1 Initial Temperature and Quarkonia Yields

The measurement of the charmonium (J/ψ , ψ') and bottomonium (Υ , Υ' , Υ'') resonances through the process $Pb + Pb \rightarrow Q\bar{Q} + X \rightarrow \mu^+\mu^- + X$ at $\sqrt{s_{NN}} = 5.5$ TeV provides crucial information about the many-body dynamics of high-density QCD matter. Because of the enhanced yield of charm quarks, the formation of J/ψ by recombination may also become significant. In addition, heavy-quark production proceeds mainly through gluon-gluon interactions and, as such, is sensitive to saturation effects on the low- x nuclear gluon density. Any discrepancies between the measurements in Pb+Pb collisions and the expectations based on p+p quarkonia cross sections will thus provide valuable information not only on the thermodynamic properties of the produced partonic medium, but also on the initial-state modifications of the nuclear gluon distribution.

Sequential suppression of heavy quarkonia production is generally agreed to be one of the most direct probes of quark gluon plasma formation. Lattice QCD calculations of the heavy-quark potential indicate that color screening dissolves the ground-state charmonium and bottomonium states, J/ψ and Υ , at $T_{\text{diss}} \approx 2T_c$ and $4T_c$, respectively. While charmonium has been studied in heavy ion collisions at the SPS and at RHIC, the clarification of some important remaining questions requires equivalent studies of the Υ family, which is only possible at LHC energies. In particular, studies of Υ and Υ' production as a function of p_T have been predicted to be sensitive to the temperature of the dense medium, directly addressing the first of the science questions listed above [3]. CMS has the best detection capabilities for the Υ family in terms of acceptance, resolution, and statistical power of any existing or planned heavy ion detector and is uniquely poised to make this crucial measurement.

2.2.2 Transport Properties of the Medium and Jet Physics

Studies of the propagation of partons through the matter formed in heavy ion collisions provide access to the transport properties of the dense medium, one of the fundamental science questions. Measurements of leading hadron production and correlations at RHIC, described in more detail below, have already pointed to strong collective effects governing high p_T phenomena. The new kinematic regime at LHC and the new probes available through precision vertexing and large coverage calorimetry will enable decisive studies of the medium properties.

Within the framework of QCD, hard collisions between two incoming hadrons can be described, at leading order, as collisions between single partons. The large Q^2 of these partons causes them to materialize immediately after the collision. Through their interactions with the evolving medium, they provide a unique probe of the transport properties of the medium [4, 5]. However, the interpretation of the existing data suffers from the fact the interacting partons are only indirectly measured via the observation of leading hadrons. Measurements of fully formed jets above the background of soft hadron production require transverse energies of $E_T > 50$ GeV/c, outside the range accessible at RHIC. At the LHC, the production rates for such processes are several orders of magnitude larger than at RHIC, allowing high statistics systematic studies in a clean kinematic regime.

In e.g. p+p collisions, jet pairs from hard parton-parton scattering are approximately produced back to back in azimuth. In Pb+Pb collisions, one or both of the parent partons of the jets can be modified by the medium, leading to a reduction in energy and change in direction. In fact, some partons may lose so much of their initial energy that they no longer appear as individual jets. The appearance of monojets was suggested as an early measure of ‘jet quenching’ due to energy loss [6, 7]. At RHIC, evidence for such an effect has been measured through leading hadron correlations [8]. This jet suppression is absent in d+Au collisions at RHIC, strongly suggesting that the suppression is a final-state effect arising from the dense medium. Leading hadron and azimuthal correlations measurements at the LHC will include fully reconstructed jets rather than simply the leading hadrons, and will extend to significantly higher p_T .

The expected statistics for dijet production will be large enough to study dijet rates as a function of impact parameter and jet transverse energy. The suppression of dijet rates due to energy loss by hard partons is expected to be much stronger for very central collisions than for peripheral collisions. The dijet rate in A+A relative to p+p collisions can be studied by introducing a reference process, such as Drell-Yan or Z^0 production, which is unaffected by energy loss and which has a rate proportional to the number of nucleon-nucleon collisions. This normalization is helpful to remove systematic errors in the centrality determination. Within a single data set, the dijet rate can be studied as a function of collision centrality and reaction plane, providing a self-calibrating measurement of nuclear modification effects by varying the path length of medium traversed.

The dijet probe is nontrivial because the original energies of both jets are, in principle, unknown. Quark jets of known energy can be produced in reactions such as $gq \rightarrow q\gamma$ [9] or $gq \rightarrow qZ^0$ [10]. In these cases, the energy can be determined since, to tree level, the parton is produced with transverse momentum equal and opposite to that of the gauge boson which is unaffected by the presence of the medium. While the reaction $gq \rightarrow qZ^0$ is a small contribution to the total Z^0 yield [11], it is a more distinctive signature since the Z^0 is free from the high background of hadronic decays contributing to the direct photon spectrum [12]. The dependence of the energy loss on the initial parton energy or distance traversed can be studied by varying the energy of the tagged photon or Z^0 in collisions with different centrality. Since the dominant channel for high p_T γ + jet production at the LHC energies is $gq \rightarrow q\gamma$, the bulk of the detected jets originate from quarks. Given the known initial parton energy measured by the photon, it is possible to perform energy-calibrated studies of the properties of opposite-side quark

jets. This is the cleanest possible channel to study quark-jet fragmentation functions.

Quark energy loss may also be studied in this channel by measuring the distribution of transverse momentum differences between the Z^0/γ and the jet. Jet energy loss should result in an asymmetric shape of this distribution. The estimated statistics are rather low for the $Z^0 + \text{jet}$ channel. On the other hand, $\gamma + \text{jet}$ production has a large background from jet + jet production if one jet is misidentified as a photon (the leading π^0), so careful corrections need to be done. Using the high granularity hadronic and electromagnetic calorimeters of the CMS detector system, a good π^0 rejection factor is achieved for the γ signal by using cuts on the cluster shape.

Because of the expected difference in the quenching mechanism for heavy quarks, the study of b -quark jets will be another important tool sensitive to the energy loss mechanism. For heavy quarks, gluon bremsstrahlung at small angles is predicted to be suppressed. This so-called *dead cone* effect leads to a considerably smaller energy loss for heavy quarks compared to light quarks. Experimentally, heavy quark jets are tagged by reconstructing secondary vertices of the leading D or B mesons. In CMS, B meson decays can be tagged either in the semi-leptonic decay channel by looking for high transverse momentum muons with displaced vertices or by reconstructing high multiplicity secondary vertices in the hadronic decay channel. A cut on the lifetime of the secondary decay will be used to vary the relative contribution of charm or bottom quark jets in the sample. The high precision muon and charged particle tracking of CMS will provide good tagging efficiency with low contamination of light quark/gluon jets.

Related measurements are single B and D meson yields. Due to the higher collision rates and more sophisticated triggering needed for this data, these analyses likely represent an effort that would come into fruition once the CMS detector and the LHC itself have both reached a more mature state of high-luminosity operation for heavy ion running.

In summary, studies of hard parton-parton scattering via jets will provide a key tool at LHC to study the transport properties of the QCD medium. The much larger cross-sections at the LHC will provide not only better statistics than at RHIC, but allow qualitatively new observables, including fully formed jets of known energy and originating from identified partons. CMS is ideally suited for these studies using the large acceptance, high resolution calorimetry and high precision tracking. Qualitatively new measurements will include spectra of charged hadrons to more than 100 GeV/c in p_T , jet E_T spectra to several 100 GeV, studies of jet-jet correlations, studies of jets tagged as coming from heavy flavors and calibrated jets tagged via real or virtual photons or Z^0 's. As an example of the accessible science, the jet or leading hadron measurements will immediately distinguish between different descriptions of the parton energy loss, which are indistinguishable at RHIC, but diverge at high p_T at LHC.

2.2.3 Thermalization and Elliptic Flow

One of the key questions for understanding the connection between heavy ion collisions and equilibrated QCD matter as described in lattice QCD calculations concerns the approach to thermal equilibrium in the early stages of heavy ion collisions. At RHIC, studies of the elliptic flow, v_2 , of produced hadrons via their azimuthal distributions relative to the reaction plane have become the main experimental tool addressing this question. Comparison to hydrodynamic calculations suggest that at the highest energies, except for the most peripheral collisions, approximate thermal equilibrium is achieved and that, correspondingly, the produced medium is characterized by a very small shear viscosity.

Theoretical efforts to understand how equilibration is achieved and to quantify the connection of medium properties like the viscosity to the experimental observables are underway. Measurements at the LHC will provide crucial new information to the existing studies through the measurement of flow at significantly higher initial densities. This is particularly important since elliptic flow data so far exhibit a steady rise in v_2 continuing up to the highest RHIC energies.

Promising studies of flow for heavy quark flavors have appeared recently at RHIC, although with limited statistical significance so far. CMS will be able to perform these measurements with high precision. The highly segmented, large acceptance calorimeters will allow a very accurate determination of the reaction plane for each event. Measurements sensitive to heavy quark flavors, e.g. based on single muons not originating from the main vertex, will be performed over a large rapidity range and out to higher p_T than accessible at RHIC. In addition, the full suite of observables like charged hadrons, π^0 s, and jets can be studied as a function of both reaction plane and centrality,

over a wide kinematic range.

In summary, flow measurements using CMS will extend previous studies of elliptic flow in terms of observables, statistics, and p_T and rapidity range. Qualitatively new information will arise from probing the effect of the increase in initial density, thereby providing clear tests of our understanding of the approach to equilibrium and the properties of the QCD medium.

2.2.4 The Initial State and Global Observables

An accelerated nucleus may be envisioned as a lattice of valence quarks surrounded by sea quarks and gluon fields. Although these sea quarks and gluons carry only a small fraction, x , of the total nucleon momentum, their density is very high, especially for the gluons. As the energy of the ion beam is increased, the lowest x values probed decreases from a typical value at midrapidity of $x \sim 0.02$ at RHIC to $x \sim 6.7 \times 10^{-4}$ at the LHC. Expanding the rapidity coverage to the forward region further reduces the accessible x values. In the regime $1.5 \leq Q^2 \leq 10 \text{ GeV}^2$ and $10^{-5} \leq x \leq 5 \times 10^{-3}$, the gluon wavelengths are long enough that the gluons begin to overlap and interact. The resultant gluon saturation is the basis of the proposed Color Glass Condensate state.

The numbers given here are relevant for p+p collisions where the gluons in a single nucleon begin to overlap each other. However, since the nonlinear growth of the gluon density depends predominantly on the number of gluons within a particular *transverse* dimension, these effects, as well as the subsequent saturation physics, may be expected to set in at higher x for nuclei than for free nucleons. CMS has the η and p_T coverage to fully define the boundaries of the saturation regime, leading to a potentially definitive test of our understanding of these effects.

Measurements at RHIC suggest that the initial state described by the concept of parton saturation is directly reflected in the multiplicity of produced hadrons and their phase space distribution. These global features of multi-particle production, measured by PHOBOS and others, exhibit great simplicity, such as factorization into separate dependencies on energy and geometry, and limiting fragmentation scaling over a large fraction of the rapidity range. In CMS, the high tracking efficiency and low rate of fake tracks provide a precise measurement of global event characteristics. Forward angle coverage is essential to studying the complete longitudinal distribution as well as pushing to the lowest possible x values. An extrapolation of existing results for the energy dependence of particle production and flow predicts a much lower multiplicity and greater elliptic flow at the LHC than can be accommodated by most current models. In addition, many of the scaling features observed at RHIC must be strongly broken at LHC energies. Confirmation of, or disagreement with, these expectations will undoubtedly severely constrain our understanding of the initial conditions and early evolution of the dense matter formed in heavy ion collisions.

2.2.5 Event Characterization

Experience at RHIC has demonstrated that global variables, such as the charged-particle multiplicity, dN_{ch}/dy , transverse energy (both E_T and dE_T/dy), azimuthal anisotropy, and energy of neutral spectators are essential for event categorization in various analyses and for placing important constraints on fundamental properties of particle production.

The large coverage of the CMS tracking detectors and the nearly full coverage of the CMS calorimeters provides access to all these measurements with high precision. Particularly the large acceptance will allow detailed studies of trigger and event selection biases, reaction plane resolutions and non-flow correlations.

The resulting event characterization observables will be available on an event-by-event basis for all other studies, allowing measurements as a function of reaction plane, as a function of the volume of the collision region, as well as the number of nucleon participants and nucleon-nucleon collisions.

In addition, it is possible to select rare, exotic events in which unusually large numbers of particles are involved or large transverse energies are produced. CMS has the capability to both trigger on these rare processes and correlate the behavior of a variety of observables over a large acceptance region.

2.2.6 Parton Distribution Functions

Using proton-proton collisions at the highest energies, we can study the nonlinear evolution regime and probe the onset of saturation in the proton. Studies of p+Pb or d+Pb collisions at 8.8 and 6.2 TeV/nucleon, respectively, can elucidate the difference between the saturation regimes of the proton and the nucleus, providing the baseline nuclear parton distributions necessary to fully understand Pb+Pb collisions at 5.5 TeV/nucleon. It is known from low energy deep-inelastic scattering that the nuclear parton distributions are modified relative to those of the free proton. However, the measurements are not very sensitive to the gluon distributions and the low x measurements are limited to $Q^2 \ll 1 \text{ GeV}^2$. At the LHC however, the low x region can be effectively probed in the perturbative regime, up to $Q^2 \geq m_Z^2$. For example, for a 10 GeV jet produced at mid-rapidity, $x \approx 10^{-3}$. Massive gauge boson (W^\pm and Z^0) production proceeds predominantly through the $q\bar{q}$ channel rather than by gluons. Thus gauge boson production is an efficient probe of the high Q^2 quark and antiquark distributions, but will be statistically significant only at LHC energies [11].

On the other hand, direct photon production is dominated by $qg \rightarrow q\gamma$ and can efficiently study the nuclear gluon distributions. Open charm production is another gluon-dominated process. Ultra-peripheral heavy ion collisions can also be used to probe gluon distributions through the exclusive production of heavy vector mesons and photoproduction of heavy quark-antiquark pairs and jets. The effects of modified parton density evolution should manifest themselves most strongly at low Q^2 , another reason why charm, with its relatively low mass, is an important probe [19]. Since modifications of the parton densities in nuclei have been observed when neither x nor Q^2 is small, measurements of the nuclear parton distributions at the LHC can help disentangle the different regimes of parton evolution. Clearly, at sufficiently high mass and p_T scales, the evolution of the proton and nuclear parton densities will again be governed by normal Q^2 evolution. The near hermeticity of CMS makes it ideal for these essential studies.

2.2.7 Forward Physics

The forward detectors, including CASTOR and the ZDCs, will enhance the entire CMS physics program. Forward coverage is essential for measurements of the parton distribution functions, particularly the gluon distributions, in protons and nuclei. CMS will be able to study momentum fractions as low as $x \sim 10^{-6} - 10^{-7}$ at scales of a few GeV^2 in p+p collisions, far exceeding the reach of previous parton density measurements. Thus, the nonlinear evolution of the parton densities can be fully mapped out, thereby illuminating the nature of the predicted gluon saturation, denoted the Color Glass Condensate, CGC [20].

At RHIC, both PHOBOS and BRAHMS have shown the value of studying Au+Au and d+Au collisions over a wide pseudorapidity region. Near beam rapidity, there exists a region of extended longitudinal scaling (also called limiting fragmentation) that increases with $\sqrt{s_{NN}}$ [21, 22]. This effect has been observed in particle ratios, integrated elliptic and directed flow, differential elliptic flow, radial flow of photons, and quark chemical potentials [23, 24, 25, 26, 27, 28]. Both the mid and forward rapidity results suggest that there are several strong constraints on the system that are set very early on and, thus, may be strongly influenced by the initial state of the nuclei before they collide. These observations may be, in part, the result of initial state saturation effects.

In heavy ion collisions, measurements in the forward region can be used to study baryon density effects on particle production, essentially changing the chemistry of the quark-gluon system [29, 30]. Thermal and chemical analysis of the current data in different rapidity slices suggests that the system has less transverse flow and fewer degrees of freedom at forward rapidity [31]. Based on extrapolations from BRAHMS d+Au data, we expect CASTOR to cover the region of maximum baryon density at the LHC [29], thus allowing a study of partonic matter over the widest possible range of baryo-chemical potential.

Finally, a detector with as close as possible to fully hermetic coverage, including forward angles, is important for the study of diffractive processes, a significant fraction (10-20%) of the total cross section at high energies. Diffractive events are characterized by large rapidity gaps between the collision products. As such, these events, which can be studied in p+p, p+A and ultra-peripheral A+A collisions, are very clean and distinctive. Hard diffractive production of heavy quarks and jets can be used to quantify the structure of the Pomeron [32].

2.2.8 Proton-Nucleus Collisions

Proton-nucleus collisions provide the cleanest measure of the initial state for nucleus-nucleus collisions and are also very interesting in their own right. It has been suggested that, when viewed by a fast probe, a heavy nucleus may resemble a sheet of highly correlated gluons, the CGC. In such a state, two soft gluons can fuse to produce one harder one. This effect results in a suppression of low p_T hadrons in d+Au collisions compared to p+p. Such effects should be stronger at forward rapidities where the momentum fraction of the partons in the nucleus is lower and their density is higher. Indeed there is some evidence from BRAHMS that such effects have already been seen in d+Au Collisions at RHIC [33]. CMS can tune the wavelength of the probe by looking at fragments that emerge over a wide range of rapidities. For central rapidities, $|\eta| < 2.4$, the full suite of silicon tracking, photon, jet, and muon measurements give a very detailed description of the collisions. For forward rapidities, CMS will certainly reconstruct jets up to $|\eta| < 5$. The combination of CASTOR and the T2 multiplicity detector may also allow us to extend jet reconstruction up to $|\eta|= 7$. Another interesting prediction of color glass and saturation models is the existence of monojets. For parton collisions below the saturation scale, momentum is conserved by the reaction of the whole condensate rather than by a single opposite-side parton.

The almost hermetic CMS calorimeter coverage will make very precise measurements of energy stopping possible in p+A collisions. These data will provide strong constraints on the initial conditions in A+A collisions. As a side benefit, p+A measurements will serve to “calibrate” the energy scale of ultra-high energy cosmic ray experiments. Currently, these experiments rely on extrapolations of GeV-scale measurements to the PeV regime in order to simulate their detector response. RHIC data has already started to influence these Monte Carlo codes but p+A data from the forward detectors of CMS would provide extremely useful input to ultra-high energy cosmic ray studies.

Peripheral p+A and A+A collisions at RHIC and LHC produce extremely large photon fluxes. The RHIC ZDCs have been used to study these events [34]. At the LHC, the maximum photon energy is 160 GeV compared with 3 GeV at RHIC. This higher energy will allow production of quarkonia and jets in photon-photon and photon-nucleus collisions, opening up the possibility of precision QCD measurements down to $x \sim 10^{-4}$ for Υ production. Photoproduction is another way to study gluon saturation and the ZDCs will be important for triggering on such events. Members of the US CMS heavy ion group have been involved in studies of high p_T jet photoproduction in p+p and p+A collisions at FNAL [35].

2.2.9 Ultra-Peripheral Collisions

Ultra-Peripheral Collisions (UPCs) can address a number of physics topics ranging from nuclear parton distributions to meson spectroscopy. Gluon distribution functions can be measured at heavy ion colliders by studying photoproduction ($\gamma g \rightarrow Q\bar{Q}$ at leading order) of heavy quarks, primarily charm and bottom [36] while top is also observable at the LHC [37]. RHIC is expected to extend the previous measurements to higher Q^2 while the LHC will be able to reach much smaller x values at high Q^2 . Gluon saturation effects may also have an influence here. For example, in a CGC [38], heavy quark photoproduction at the LHC would be reduced relative to predictions using the parton distributions of a free proton.

The hermeticity of CMS is an advantage for studies of ultra-peripheral collisions relative to other experiments at the LHC. The largest possible rapidity coverage is important to reduce backgrounds. The ZDC described in Appendix D is an essential component of ultra-peripheral studies because the number of neutrons following the incoming beam trajectory reflects the nuclear state after the collision.

2.2.10 Proton-Proton Physics

From a heavy ion perspective, more elementary proton-proton collisions serve three important functions: as a baseline for what to expect for global properties such as particle yields, as a source of reference data to investigate the effect of the nuclear medium on observables such as transverse momentum distributions, and as a way to study fundamental features such as the dependence of particle production on available energy. All three of these roles have been amply demonstrated at RHIC. It is clear that a large part of our understanding of heavy ions at the LHC will also depend on high quality p+p data. Because of the large jump in proton-proton center of mass energy, there is also a good chance that surprises may await in such basic global properties as total multiplicity, distributions in pseudorapidity and transverse momentum, two-particle correlations in rapidity and azimuth, and the dependence

of various observables on multiplicity. The broad coverage of the CMS detector, in particular the forward coverage provided by CASTOR and the ZDC, will allow unique measurements in the areas of low- x and diffractive physics. These studies are particularly relevant since the parton saturation concepts developed in response to existing RHIC and HERA data should be even more influential in the LHC energy regime. Due to the limited bandwidth and the focus on very rare processes, observables such as open and hidden charm and bottom quark production, low p_T jets, and the more global properties mentioned above will be secondary priorities in the high energy program. By developing the triggering and analysis tools necessary for these measurements which are critical to heavy ion physics, an opportunity exists to exploit similar studies in elementary collisions.

The first pilot p+p run will take place in late 2007, currently planned to be at a center of mass energy of 900 GeV. The initial low luminosity p+p run in 2008 will be at a center of mass energy of 14 TeV, higher than the planned Pb+Pb energy of $\sqrt{s_{NN}} = 5.5$ TeV. However, since many properties are known to scale roughly logarithmically with \sqrt{s} , it should be possible to interpolate reasonably accurately to the heavy ion beam energies using LHC and Tevatron data. These early p+p results will be critical for interpreting the first Pb+Pb data expected in late 2008 or early 2009. Reference p+p data at the lower energy must be taken as part of the one month per year allocated to heavy ion running and, therefore, will most likely not occur until later in the program. The significantly lower expected collision rate in the first p+p run at 14 TeV will allow a cleaner environment for determining global observables and other properties of interest to the heavy ion program. Also, the DAQ bandwidth at the higher luminosities will be more fully saturated with triggers on the rarer processes of more interest to the high energy community.

Active participation in the 2007 p+p run has important practical consequences. Due to the relatively short time allocated to Pb beams, it is critical that the trigger hardware and software, as well as the off-line analysis tools are fully installed, commissioned, and tested well in advance. The triggering constraints and the most important calibration and alignment parameters will certainly have both overlap and differences between high energy and heavy ion analyses. Therefore, work by the heavy ion group on the early p+p data will maximize the efficiency of the following Pb+Pb run as well as enhancing the interesting physics potential of this data from a broader perspective.

Process	Pb+Pb $\sqrt{s_{NN}} = 5.5$ TeV $\mathcal{L} = 5 \times 10^{26} \text{ cm}^{-2} \text{ s}^{-1}$		p+Pb $\sqrt{s_{NN}} = 8.8$ TeV $\mathcal{L} = 1.4 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$	
	Yield/ 10^6 s	Ref.	Yield/ 10^6 s	Ref.
$ \eta \leq 2.4$				
jet($p_T > 50$ GeV)	2.2×10^7	[15]	1.5×10^{10}	[14]
jet($p_T > 250$ GeV)	2.2×10^3	[15]	5.2×10^6	[14]
Z^0	3.2×10^5	[11]	6.8×10^6	[14]
W^+	5.0×10^5	[11]	1.1×10^7	[14]
W^-	5.3×10^5	[11]	1.1×10^7	[14]
all phase space				
$c\bar{c}$	9.0×10^{10}	[16]	2.0×10^{12}	[16]
$b\bar{b}$	3.6×10^9	[16]	8.2×10^{10}	[16]
$J/\psi \rightarrow \mu^+ \mu^-$	2.4×10^7	[17]	5.5×10^8	[17]
$\Upsilon \rightarrow \mu^+ \mu^-$	1.5×10^5	[17]	3.5×10^6	[17]
$\Upsilon' \rightarrow \mu^+ \mu^-$	3.7×10^4	[17]	8.4×10^5	[17]
$\Upsilon'' \rightarrow \mu^+ \mu^-$	2.2×10^4	[17]	5.2×10^5	[17]

Table 1: The yield of hard probes in a 10^6 s LHC run.

2.2.11 Expected Rates for Hard Probes

Semi-hard and hard parton processes will be produced abundantly at the LHC [13]. For the first time in heavy ion collisions, high- p_T jets and photons, quarkonia, and W^\pm or Z^0 bosons, characterized by large squared momentum transfer, $Q^2 > 50 \text{ GeV}^2$, will be produced with usable statistics. As discussed above, these probes open a broad range of new opportunities for investigating the properties of the dense medium formed in heavy ion collisions.

To better illustrate the high rates available for hard probes at the LHC, Table 1 presents the minimum bias jet and gauge boson rates in the region $|\eta| \leq 2.4$ as well as the total heavy flavor and quarkonium rates calculated in perturbative QCD in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.5$ TeV and p+Pb collisions at $\sqrt{s_{NN}} = 8.8$ TeV [11, 14, 15, 16, 17]. The results are given for a 10^6 second LHC run and assume a luminosity of $5 \times 10^{26} \text{ cm}^{-2}\text{s}^{-1}$ for Pb+Pb [18] and a maximum p+Pb luminosity of $1.4 \times 10^{30} \text{ cm}^{-2}\text{s}^{-1}$ [14]. Conventional nuclear shadowing, the only nuclear effect included, is typically rather small for jets and gauge bosons but can be large for heavy quarks and quarkonium. The quarkonium rates include the branching ratios to lepton pairs.

2.3 Physics Interests of the US groups

The members of the US CMS heavy ion group have a broad range of experience and expertise in heavy ion physics as well as other areas. The current specific interests of the individual groups are listed in Table 2. The “Global” area includes, among other overall properties, the dependences on energy, centrality, and pseudorapidity of particle yields as well as correlations, fluctuations, and flow. The subject of jet suppression, believed to result from partonic energy loss in the dense medium, is covered in “High p_T ”. These studies include unidentified charged particles found in the tracker, jets found in the calorimeters, and π^0 s reconstructed from photons in the electromagnetic calorimeters. The “Tagged jets” category includes more detailed investigation of energy loss made using heavy quark jets tagged by a displaced vertex, as well as jets with a less highly interacting partner such as a Z^0 . While requiring higher luminosity for meaningful statistics, these data will probe more detailed characteristics of partonic energy loss. The results will also quantify open charm and bottom quark yields which are critical to the interpretation of the next topic in the table. The “Quarkonia” category encompasses the suppression of heavy quark mesons due to disassociation in the dense partonic medium. In addition, possible J/ψ enhancement due to quark coalescence can be investigated. The “Forward” physics area includes measurements of particles and jets at extreme pseudorapidity as well as the use of forward detectors to study and trigger on ultraperipheral collisions. As discussed in Section 2.2.10, all of these physics topics will also be investigated in p+p data, both as a reference for the heavy ion results and as a more general study of the properties of particle production.

These various physics topics require a wide range of event statistics as well as detector and triggering capabilities. With a range of physics interests, the US CMS heavy ion group is well positioned to exploit exciting physics possibilities from the earliest low luminosity runs up through the later high statistics data. Early results could easily redirect the focus of later studies. Through interactions within and between of the various groups, the full breadth of the physics will be well covered, allowing any overlap between the various analyses to be fully explored.

Group	Global	High p_T	Tagged jets	Quarkonia	Forward
UCDavis	✓			✓	
UIC	✓		✓		✓
Iowa	✓	✓			✓
Kansas	✓	✓			✓
LANL			✓		
Maryland	✓				✓
MIT	✓	✓	✓		
Minnesota	✓		✓		
Rice		✓			✓
Vanderbilt	✓		✓	✓	

Table 2: Current interests of the US CMS heavy ion groups in several physics areas. See text for discussion.

2.4 Physics Schedule

The timescale with which various physics topics can be addressed is determined by the expected beam time and luminosity allocated to the different modes of running at the LHC. Table 3 shows the current plan for the LHC along with the expected number of events useful for the heavy ion physics program. For p+p collisions, only the expected number of minimum bias events is counted since the highly restrictive triggers focus to a large degree on observables of interest only to the high energy program.

Calendar Year	Beam	Million Events
2007	Pilot p+p	48
2008	Physics p+p, low luminosity	110
	Pilot HI at 1/20 luminosity	60
2009	Physics p+p, full luminosity	68
	Physics HI, full luminosity	120
2010	Physics p+p	42
	Physics HI	120
2011	Physics p+p	42
	Physics HI	180
2012	Physics p+p	30
	Physics HI	120
	p+p at HI beam energy	250

Table 3: The projected run schedule at the LHC. Note that only the minimum bias events of interest to the heavy ion program are counted for the p+p runs.

Using the steady accumulation of events, a range of physics topics starting with basic global event properties and moving onto more and more rare processes can be addressed. A preliminary schedule of physics goals is shown in Table 4. In order to accomplish these goals, a clear set of experimental tasks must be accomplished.

- FY2006: Develop trigger hardware and software as well as physics analysis algorithms. Beam tests with prototype ZDC. Construct final ZDC and prepare for installation.
- FY2007: Continue trigger and physics analysis development. Data processing of min-bias p+p. Commission ZDC, HI computer center, and DAQ system.
- FY2008: Continue software development and p+p analysis. Implement HLT trigger. Upgrade HI computer center.
- FY2009: Continue software development. Optimize HLT trigger. Pb+Pb data taking and analysis. Upgrade HI computer center.
- FY2010 – FY2012: Ongoing software development, HLT trigger optimization, and Pb+Pb data taking. HI computer center maintenance.

Calendar Year	Physics Goals
2006	Monte Carlo simulations: <ul style="list-style-type: none"> • Capabilities of the detector for HI • Development of reconstruction algorithms • ZDC test beam calibration
2007	Monte Carlo simulations: <ul style="list-style-type: none"> • Capabilities of the detector for HI • Trigger performance for HI 900 GeV p+p data: <ul style="list-style-type: none"> • First opportunity to test detector and trigger performance • p+p charged multiplicity • p+p particle and jet spectra • ZDC performance
2008	Reference p+p data: <ul style="list-style-type: none"> • p+p charged multiplicity • p+p particle and jet spectra • p+p jet studies: dijet, jet-γ, jet-Z^0 • p+p quarkonium and heavy quark measurements First HI data: <ul style="list-style-type: none"> • Charged particle multiplicity and flow • Spectra of charged particles to $p_T < 100$ GeV • Observation of jets with $E_T < 200$ GeV
2009	Increased statistics reference p+p data: Increased statistics HI data: <ul style="list-style-type: none"> • Charged multiplicity and flow vs centrality • Charged particle spectra to $p_T < \text{multi-}100$ GeV • Jet studies for jets with $E_T < \text{multi-}100$ GeV • Initial open c and b quark studies
2010	Further statistics for p+p and HI data: <ul style="list-style-type: none"> • Detailed jet fragmentation studies, multi-jets • First J/ψ, Υ family results • First jet-γ and jet-Z^0 observations • Detailed open c and b quark studies
2011	Further statistics for p+p and HI data: <ul style="list-style-type: none"> • Detailed multi-jet studies • Detailed J/ψ, Υ family, centrality dependence • Jet-γ, jet-Z^0 studies, centrality dependence • Extensive heavy quark studies
2012	p+p reference at 5.5 TeV: Refine analysis with improved p+p reference Increased statistics HI data: <ul style="list-style-type: none"> • Detailed studies of all rare channels, centrality dependence

Table 4: A preliminary schedule for the physics goals of the heavy ion program at the LHC.

3 The CMS Detector

The capabilities of the CMS detector provide a unique opportunity to exploit the new physics potential provided by heavy ion collisions at the LHC. The connection between detector performance and specific physics topics is discussed in detail in Section 4. This section provides an introduction to the most important hardware aspects of the experiment.

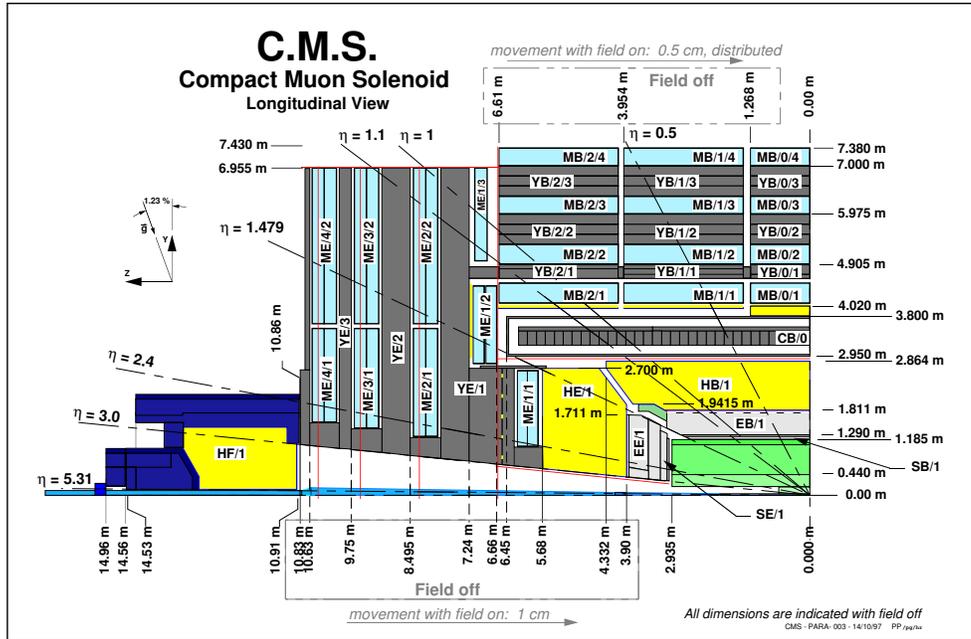


Figure 1: A vertical slice through one quadrant of the CMS detector. The other three quadrants are identical and symmetric about the center of the interaction region and about the beam axis.

The CMS detector was designed to provide tracking and calorimetry with high resolution and granularity over the full azimuthal angle as well as a very large range in rapidity. The primary emphasis is on precisely detecting muons, electrons, photons, and jets, but significant capability exists for other products of heavy ion reactions. The various elements, either singly or in combination, can be used to perform identification of a wide array of particle species. The electronics and data acquisition systems allow a very fast initial readout as well as sophisticated multi-level triggering. More detailed descriptions are given in the Technical Design Reports [39, 40, 41, 42].

The overall layout and dimensions of the apparatus are illustrated in Fig. 1 which shows one quadrant of a vertical slice. The full detector is symmetric about both the beam axis and the center of the nominal interaction region. A 3-dimensional cut-away view is shown on the title page. Figure 2 shows schematic representations of the response to various types of particles superimposed on a transverse slice through the detector. The dominant element of the experiment is the magnet, the coil of which is labeled “CB” in Fig. 1. The superconducting solenoidal coil is roughly 13 m long, 6 m in diameter, and provides a 4 T field throughout the inner portion of the detector. This inner region holds the silicon tracking system, arrayed throughout the unlabeled shaded region starting at location (0,0) at the bottom right in Fig. 1, as well as the electromagnetic (“EE” and “EB”) and hadronic (“HE” and “HB”) calorimeters. Additional information about particles at high pseudorapidity is provided by a calorimeter (“HF”) located near the beam line about 11 m from the interaction point. The regions outside the coil as well as in the forward and backward directions are filled with tracking (“ME” and “MB”) and absorbers (integrated into the magnet return yoke, “YE” and “YB”) for detecting muons. The central tracking covers $|\eta| < 2.5$, the central calorimeters cover $|\eta| < 3$ and the forward calorimeters extend the coverage to $3 < |\eta| < 5$. Muons can be tracked and identified inside $|\eta| < 2.4$, roughly the same region covered by the inner Si tracker.

Very far from the interaction point, and therefore not shown in the figure, the experiment includes a suite of detectors designed to study particles emitted at very high pseudorapidity. The CASTOR calorimeter [43] covers

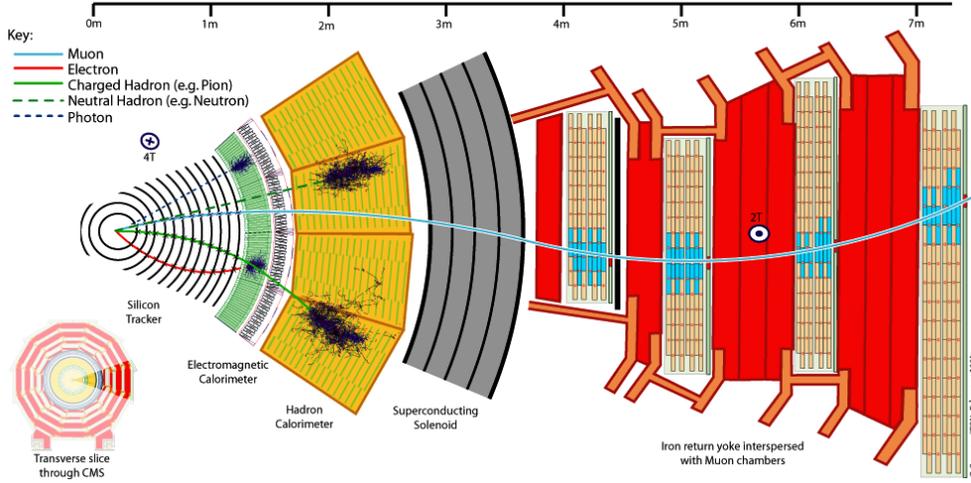


Figure 2: A transverse slice through one segment of the CMS detector.

$5 < |\eta| < 7$ and the TOTEM Roman pots [44] extend this to $7 < |\eta| < 10$. Finally, the Zero Degree Calorimeter (ZDC) sits 140 m away which is behind the first accelerator magnet. This detector is primarily sensitive to neutron spectators from the colliding ions. As one of the major hardware contributions of the US heavy ion group, this detector is described in more detail in Appendix D.

3.1 Inner Si Tracker

Data from the Si detector system provide precision tracking of particle trajectories close to the interaction point. This information can be used to determine the momentum of particles impinging on the outlying calorimeters and muon detectors. Track vectors can be combined to give a precision determination of the vertex. Finally, a high-segmentation measurement of the distribution of charged particles in azimuth, pseudorapidity, and transverse momentum can be extracted. The Si wafers are arranged in 10 cylindrical layers coaxial with the beam pipe as well as 12 disk-shaped layers perpendicular to the beam in both the forward and backward directions. It is important to note that the segmentation of the tracker is set by the requirements for momentum resolution and, as a result, the occupancies are low even for the most central Pb+Pb collisions.

Wafers closer to the beam pipe (3 layers in the barrel and 2 in the endcaps) are subdivided into pixels, while those farther away are segmented into strips. The pixel wafers in the barrel, covering pseudorapidities up to $|\eta| = 2.1$, are situated 4.5 cm, 7.5 cm and 10 cm from the beam. With individual readout areas of $100 \times 150 \mu\text{m}^2$, there are a total of more than 9.6 million, 16 million, and 22.4 million pixels in the three layers. The precision for high p_T tracks is around $15\text{-}20 \mu\text{m}$. The closest Si strip counters are arranged in 4 layers in the barrel and 3 in the endcap. The wafers have individual strips with a width of $80 \mu\text{m}$ and length of 6.1 cm. In the barrel, these inner strip detectors are mounted between 210 and 635 mm from the beam. The outer portion of the tracker is comprised of 6 barrel layers and 9 endcap disks, with strip pitch ranging from 122 to $183 \mu\text{m}$ and strip lengths of 9.1 cm.

An important distinction of the CMS readout for both the Si strips and pixels is the inclusion of pulse height information instead of simple digital counts. Experience at RHIC, especially from PHOBOS, has shown that this information is vital to correct for a variety of detector effects such as baseline drift and cross-talk. The magnitude of such effects is frequently larger in the higher multiplicity heavy ion environment. In addition, particle identification in the low transverse momentum sector using energy loss in the Si has been very successful in PHOBOS and the capability of making similar measurements using CMS is being investigated.

3.2 Calorimeters

The extensive coverage of both electromagnetic and hadronic calorimetry is one of the distinguishing features of the CMS apparatus. Similar to the Si tracker, electromagnetic and hadronic elements are arranged in barrels and

endcaps located inside the magnet coils with additional electromagnetic and hadronic elements located farther away.

The barrel electromagnetic calorimeter covers the range $|\eta| < 1.48$. The crystals have dimensions of 22 mm square (corresponding to $\Delta\eta \times \Delta\phi = 0.0175 \times 0.0175$) with their inner face at a radius of 1.29 m. Their length of 230 mm provides $25.8 X_0$. The endcap electromagnetic segment extends over $1.48 < |\eta| < 3$. Because the crystal size is uniform, there is a decreasing granularity in η closer to the beam direction. In the most forward region, the segmentation is $\Delta\eta \times \Delta\phi = 0.05 \times 0.05$. The full electromagnetic calorimeter contains almost 76000 scintillating crystals of PbWO_4 , read out using avalanche photodiodes in the barrel and vacuum phototriodes in the endcaps. A preshower detector, consisting of two layers of orthogonally-oriented silicon strip detectors located behind 1 and $2 X_0$ of lead, covers a portion ($1.65 < |\eta| < 2.6$) of the front of the endcap calorimeter.

The hadronic barrel and endcap are sampling calorimeters with alternating 4 mm thick scintillators and copper plates, the latter being 5 and 8 cm thick in the barrel and endcap, respectively. The barrel portion is about 79 cm deep, corresponding to 5.15 nuclear interaction lengths at $\eta = 0$. The scintillator tiles are read out using wavelength-shifting plastic fibers.

The forward calorimeter (including both electromagnetic and hadronic sections) covers the region $3 < |\eta| < 5$ with a segmentation of $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$. For p+p collisions, the significantly enhanced hermeticity of the detector results in a more accurate determination of missing transverse energy. For heavy ion collisions, it extends the measurement of the forward energy and will provide additional capability for centrality determination. The detector is assembled from quartz fibers (predominantly sensitive to Čerenkov light from neutral pions) embedded in a steel absorber matrix.

With such an extensive coverage for calorimetry, the nature of the data available for physics analyses is qualitatively different. Beyond simply increasing the statistics or pseudorapidity extent, the algorithms themselves can be significantly different. As one example, the effect of different cone sizes in jet finding is easily investigated. An addition of particular utility in heavy ion collisions is the possible use of information from other regions of the detector to establish the background underlying a jet.

3.3 Muon System

The muon detection system covers $|\eta| < 2.4$, including $|\eta| < 1.3$ in the barrel and $0.9 < |\eta| < 2.4$ in the endcaps. Additional resistive plate chambers (6 layers in the barrel and 4 in the endcaps) designed for trigger purposes cover $|\eta| < 2.1$. The barrel tracking is provided by drift tubes with 2.5 m long anode wires, arranged in 4 groups (called “stations”) of 12 planes each. Each station provides an independent 3-dimensional point on the muon trajectory. In the endcap, particles are detected using 4 stations, each comprised of 6 layers of cathode strip chambers. All 24 layers measure both the radial and azimuthal position of the trajectory.

3.4 Detectors in the Forward Region

The CMS coverage beyond $\eta = 5$ is unique at the LHC. The CASTOR calorimeter, T2 GEM detector, TOTEM, and ZDC extend particle detection capability to $|\eta| < 10$ for charged particles and to zero degrees for neutrons and photons. CASTOR, which covers $5 < |\eta| < 7$ contains 10.6 interaction lengths of quartz fibers and tungsten plates, subdivided into 8 segments in ϕ and 18 segments along the beam direction. The ZDC is described in detail in Appendix D.

4 CMS Detector Performance for Heavy Ion Collisions

The RHIC discoveries have not only transformed our picture of nuclear matter at extreme densities, but also greatly shifted the emphasis in the observables best suited for extracting the properties of the initial high-density QCD system [45]. Examples of these observables include the elliptic flow coefficient, very high p_T jets, and heavy quarkonia. This change of paradigm for the probes of the medium suggests the need for detectors with large acceptance, high rate capability, and high resolution, leading to a convergence of experimental techniques between heavy ion and particle physics. The proposal to use CMS for heavy ion collisions takes this development to its logical conclusion, leveraging the extensive resources that have gone into the development and construction of the detector. The utility of the convergence of heavy ion and particle physics techniques is illustrated by the extent to which CMS fulfills the criteria for an ideal heavy ion detector.

1. **High rate:** CMS is designed to deal with p+p collisions at luminosities of up to $10^{34} \text{ cm}^{-2}\text{s}^{-1}$, corresponding to p+p collision rates of 10^9 Hz. As a result, the fast detector technologies chosen for tracking (Si-pixels and strips), electromagnetic and hadronic calorimetry, and muon identification will allow CMS to be read out with a minimum bias trigger at the full expected Pb+Pb luminosity. This fast readout will allow detailed inspection of every event in the high level trigger farm. The bandwidth of the farm is sufficient to run complex analysis algorithms on each of the events, allowing a complete selection and archiving of events containing rare probes such as extremely high p_T jets or events having unusual global properties.
2. **High resolution and granularity:** At the full p+p luminosity, there will be, on average, 25 collisions per bunch crossing. To be able to disentangle very high momentum observables with $p_T > 500 \text{ GeV}/c$ in this environment, the resolution and granularity of all detector components has been pushed to the extreme, consequently making the detector ideally suited to the high multiplicity conditions in central heavy ion collisions. The high granularity of the Si-pixel layers, in combination with the 4 T magnetic field, results in the world's best momentum resolution, $\Delta p_T/p_T < 1.5\%$ up to $p_T \approx 100 \text{ GeV}/c$. At the same time, a track-pointing resolution of less than $50 \mu\text{m}$ (less than $20 \mu\text{m}$ for $p_T > 10 \text{ GeV}/c$) is achieved. At $dN_{\text{ch}}/dy \approx 3000$, tracks can be reconstructed with an efficiency of $\sim 80\%$, which is more than adequate. The electromagnetic calorimeter can be used to find jet locations with resolutions in η and ϕ of 0.028 and 0.032, respectively.
3. **Large acceptance tracking and calorimetry:** CMS includes high resolution tracking and calorimetry over 2π in azimuth and a uniquely large range in rapidity. The acceptance of the tracking detectors, calorimeters, and muon chambers can be seen in Fig. 3. The Zero Degree Calorimeters ($|\eta_{\text{neutral}}| > 8.0$) and the CASTOR detector ($5.2 < |\eta| < 6.6$) will allow measurements of low- x phenomena and particle and energy flow at very forward rapidities. The CMS and TOTEM experiments will constitute the largest acceptance system ever built at a hadron collider.
4. **Particle identification:** At the LHC, charm and bottom quarks will be copiously produced. The large acceptance, high resolution muon system, in combination with the tagging of secondary decays by the silicon tracker, will allow studies of the interaction of identified quarks with the medium. In addition, the physics of meson vs. baryon production at large p_T can be studied using the results for reconstructed π^0 s, as well as the information provided by the silicon tracker in combination with the electromagnetic and hadronic calorimeters. In the low transverse momentum regime, further studies will be performed to assess the information that can be obtained from measurements of specific ionization in the silicon detectors as well as the reconstruction of hadronic resonances using invariant mass analysis.

4.1 Tracking Performance

The CMS performance in most categories far exceeds the capabilities of existing or planned heavy ion detectors. However, many of the physics topics that can be addressed depend on the ability to reconstruct individual charged particles. Track reconstruction in the CMS silicon tracker in a very high track density environment is thus a key element to the success of the CMS heavy ion program.

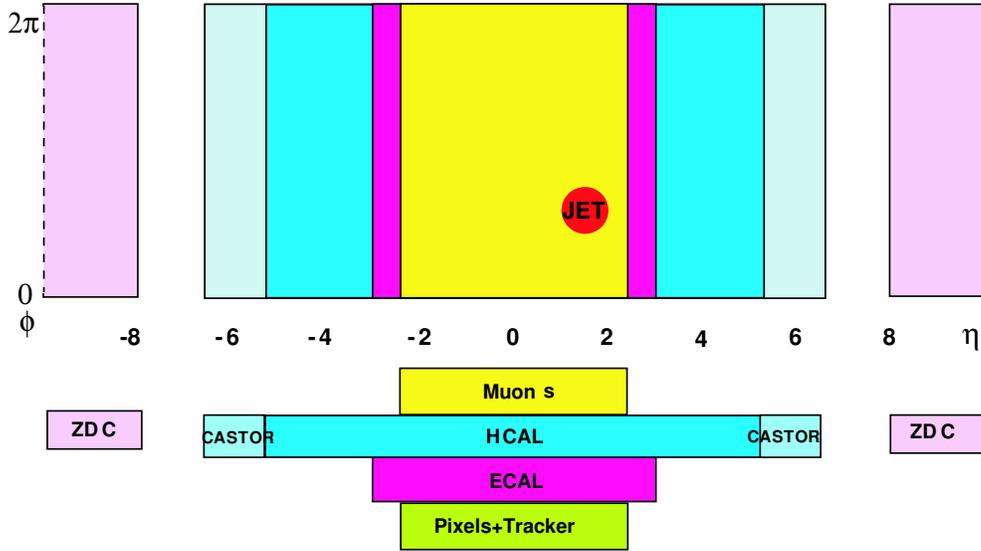


Figure 3: Acceptance of tracking, calorimetry, and muon identification in pseudorapidity and azimuth. The size of a jet with cone $R = 0.5$ is also depicted as an illustration.

The maximum charged particle density of central Pb+Pb collisions cannot easily be extrapolated from RHIC energies to the LHC. Assuming a logarithmic increase of the charged particle density with the nucleon-nucleon center-of-mass energy [46], between 1500 and 4000 charged particles per unit rapidity will be produced in central Pb+Pb collisions at the LHC. For this study the parameters of the simulation were set to produce a midrapidity charged particle density of about 3000 per unit rapidity, as shown in Fig. 4, which results in a very high detector occupancy as can be seen in Fig. 5. The combinatorial challenge resulting from the high hit density requires robust reconstruction algorithms to achieve efficient pattern recognition while maintaining a low fake rate.

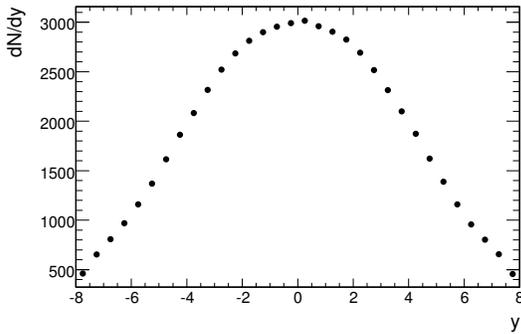


Figure 4: Simulated charged particle density in central Pb+Pb collisions as a function of rapidity. This calculation corresponds to a density close to the maximum expected from extrapolation of lower energy data.

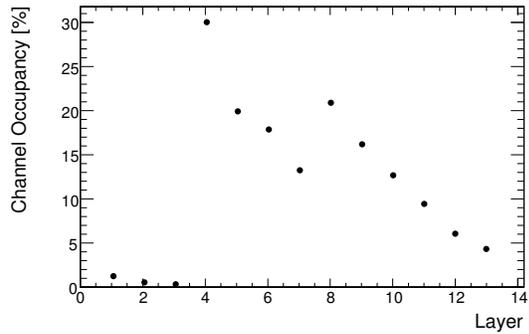


Figure 5: Channel occupancy in the barrel region as a function of detector layer for the particle yields shown in Fig. 4. Layer 1-3 correspond to the pixel detector, 4-7 to the inner and 8-13 to the outer Si-strip tracker.

The performance of the tracker in the environment of central Pb+Pb collisions was evaluated using a data sample generated with the HYDJET event generator [47]. Events simulated with this generator consist of contributions from soft particle production modeled by a hydrodynamic module and multiple hard collisions simulated with PYTHIA. Events were processed through the official CMS package for full simulation (OSCAR) and reconstruction (ORCA). The standard track reconstruction algorithm, initially tuned for p+p, was optimized for the high density environment. In order to reject fake tracks, cuts on the minimum number of required hits, a maximum

fit χ^2 , and a minimum distance of closest approach to the primary vertex (dca) were applied. Figure 6 shows the

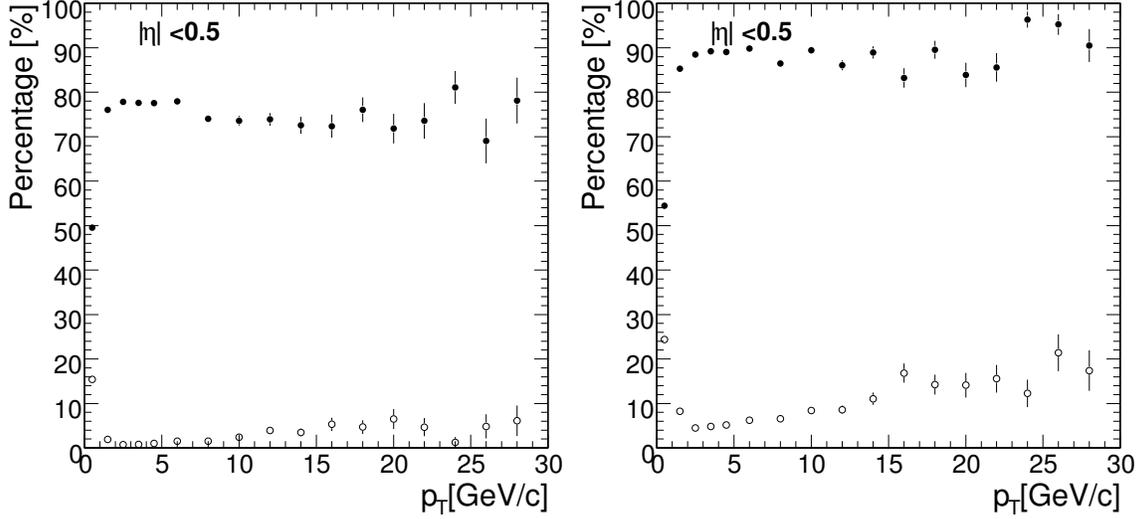


Figure 6: Track reconstruction efficiency (full symbols) and fake track rate (open symbols) as a function of transverse momentum near midrapidity for central Pb+Pb collisions with $dN_{ch}/dy \approx 3000$. (Left) Track quality cuts optimized for low fake track rate (cut on number of hits > 12 , fit probability > 0.01 , and dca < 3). (Right) Track quality cuts optimized for high efficiency (only cut on number of hits > 12).

track reconstruction efficiency and fake track rate as a function of transverse momentum in the barrel region of the tracker for two sets of quality cuts imposed in the reconstruction. The momentum and track-pointing resolution achieved in heavy ion collisions (see Fig. 7) are comparable to those in low occupancy p+p events.

Given the excellent pointing accuracy of single tracks and the high event multiplicity, the primary event vertex can be reconstructed with a resolution of better than $10 \mu\text{m}$ in all 3 dimensions. The precision of the primary event vertex reconstruction will clearly be limited only by the alignment precision of the tracker components.

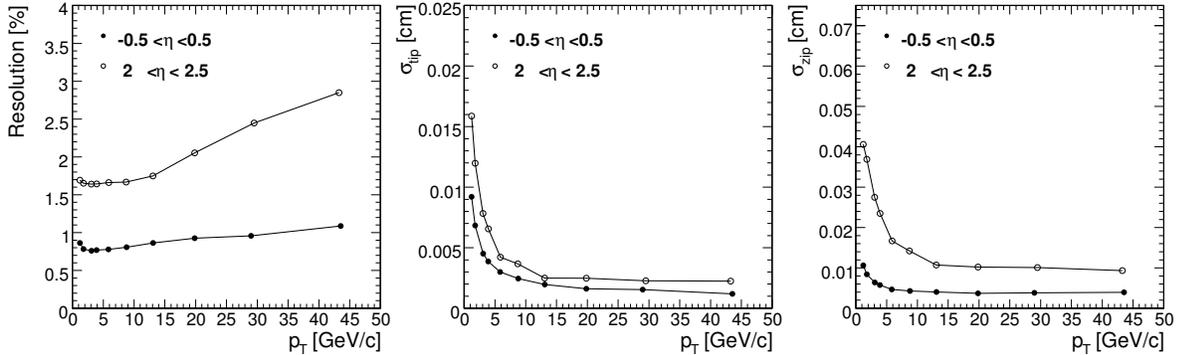


Figure 7: The p_T dependence of the track parameter resolution achieved in heavy ion events in the barrel region (full symbols) and in the forward region (open symbols). (Left) Transverse momentum resolution. (Center) Transverse track-pointing resolution. (Right) Longitudinal track-pointing resolution.

With the nominal magnetic field setting, the current tracking algorithms have been found to have good efficiency and purity down to a transverse momentum of $p_T \approx 1 \text{ GeV}/c$. For broader global measurements, preliminary feasibility studies of measuring bulk charged particle multiplicities using individual layers in the CMS silicon tracker have been performed. These early analyses will be expanded to enable a more detailed study of charged particle yields versus centrality, pseudorapidity and angle with respect to the reaction plane. In order to extend the analysis even further, there is also interest in investigating the (low) momentum tracking and particle identification

possibilities using the multiple layers of silicon detectors. The preliminary results indicate that there is a good possibility of extending the tracking limit to well below $p_T \approx 1 \text{ GeV}/c$.

4.2 Muon Reconstruction

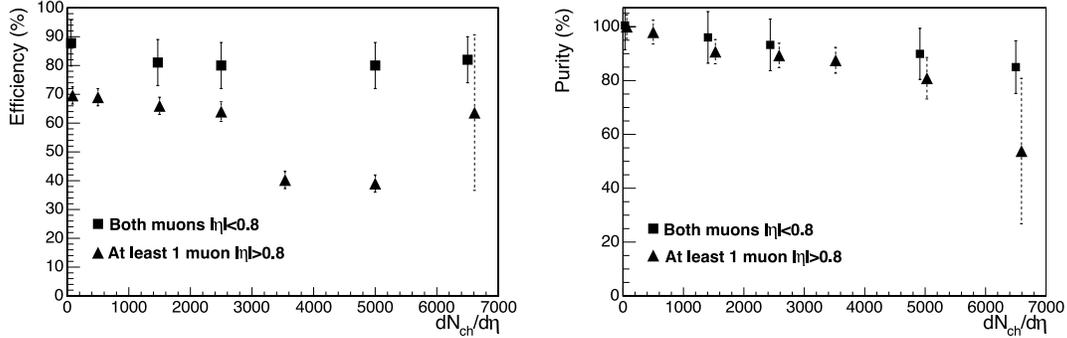


Figure 8: Υ reconstruction efficiency (left) and purity (right) as a function of the Pb+Pb charged particle rapidity density, $dN_{\text{ch}}/d\eta|_{\eta=0}$. Note that a different algorithm must be used at the higher particle densities.

CMS is an ideal detector for muon reconstruction with an acceptance spanning 4.8 units in pseudorapidity, the largest of any heavy ion detector. The proximity of the electromagnetic calorimeter to the beam axis (1.29 m) minimizes the background from pion and kaon decays. An additional background rejection of a factor of 6 (2) in the barrel (endcaps) is obtained using the excellent muon momentum resolution given by the tracking detector. For heavy ions, the muon reconstruction algorithm starts with the pixel tracking information. The inner cylinder pixel detector is placed at 4.5 cm from the beam axis and is used to determine the interaction point. Once the vertex is known, information from the muon detector and the tracking is used to reconstruct the muon momentum. Figure 8 shows the Υ efficiency and purity (where purity is defined as the ratio of true Υ reconstructed over all Υ reconstructed) as a function of charged-particle multiplicity. In the central barrel, the dimuon reconstruction efficiency is above $\sim 80\%$ for all multiplicities, whereas the purity decreases slightly with $dN_{\text{ch}}/d\eta$ but stays also above 80% even at multiplicities as high as $dN_{\text{ch}}/d\eta|_{\eta=0} = 6500$. If (at least) one of the muons is detected in the endcaps, the efficiency and purity drop due to the need for tighter track reconstruction cuts. Nonetheless, for the maximum multiplicities expected in central Pb+Pb collisions, $dN_{\text{ch}}/d\eta|_{\eta=0} \approx 2500$, the efficiency (purity) remains above 65% (90%).

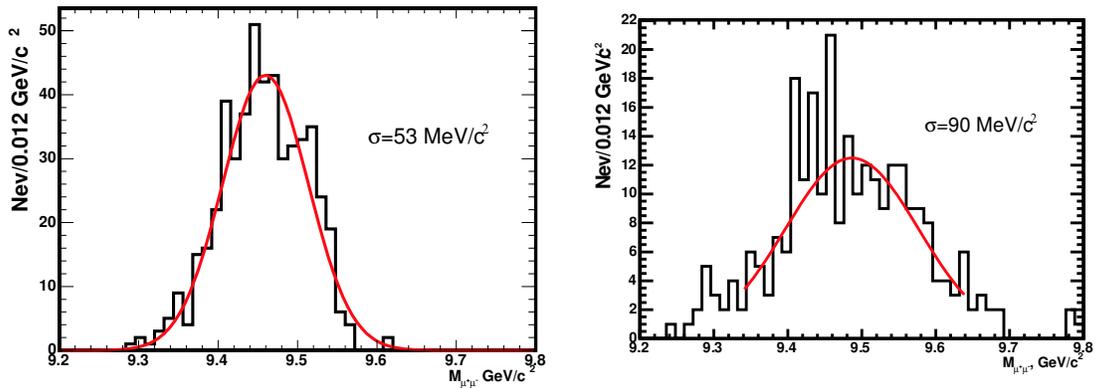


Figure 9: The reconstructed Υ mass. (Left) Without Pb+Pb background events and with both muons in $|\eta| < 0.8$ (barrel region only). (Right) In a Pb+Pb event with multiplicity $dN_{\text{ch}}/d\eta|_{\eta=0} = 2500$ and both muons in $|\eta| < 2.4$.

Figure 9 shows the mass resolution of the reconstructed muon pairs from Υ decay obtained with the full simulation and reconstruction. The good muon momentum resolution translates to an Υ mass resolution in p+p of $53 \text{ MeV}/c^2$

in the barrel pseudorapidity region, the best of all LHC detectors. When the signal is superimposed on a background event and measured in the full pseudorapidity region, the mass resolution is $90 \text{ MeV}/c^2$. This provides a clean separation between the members of the Υ family and also a significant improvement in the signal to background ratio.

To further illustrate the resolution and yield, Fig. 10 shows the dimuon mass distributions, after background subtraction, for two different scenarios: $dN_{\text{ch}}/d\eta|_{\eta=0} = 5000$, $|\eta| < 2.4$; and $dN_{\text{ch}}/d\eta|_{\eta=0} = 2500$, $|\eta| < 0.8$. Except for the ψ' , all quarkonia states are clearly visible. These yields correspond to one month of Pb+Pb running.

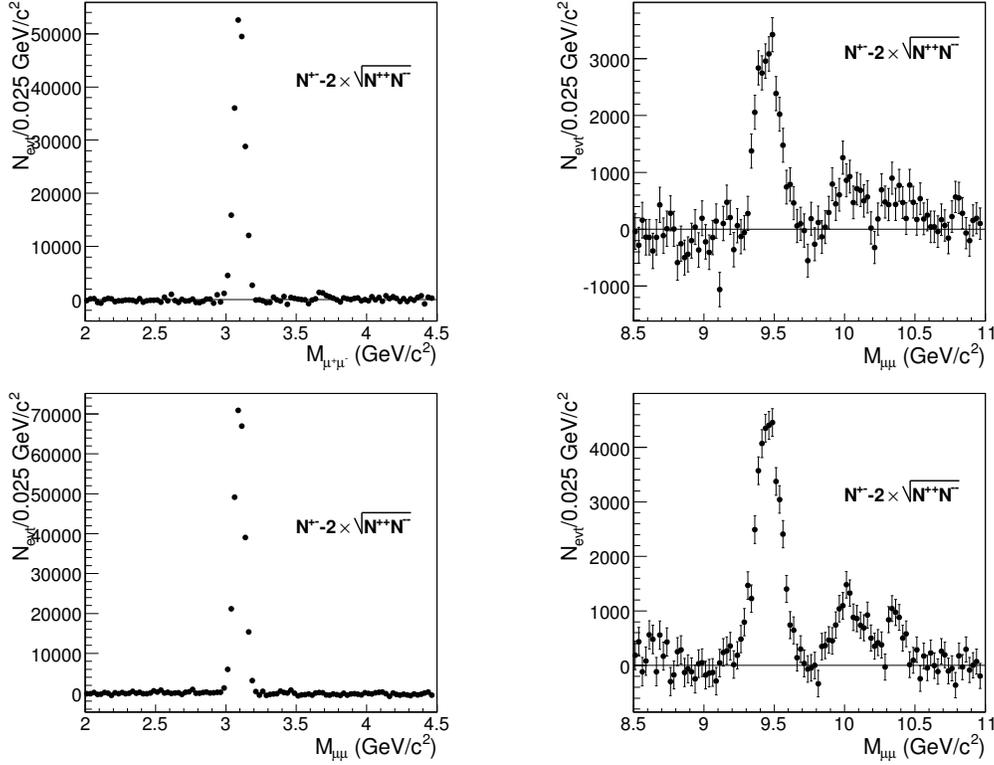


Figure 10: Signal dimuon mass distributions after background subtraction in the J/ψ (left) and Υ (right) mass regions expected after one month of Pb+Pb running. Top panels are for $dN_{\text{ch}}/d\eta|_{\eta=0} = 5000$ and $|\eta| < 2.4$; bottom panels are for $dN_{\text{ch}}/d\eta|_{\eta=0} = 2500$ and $|\eta| < 0.8$.

To date, most quarkonium studies in CMS have focused on detection through their decays to $\mu^+\mu^-$. This decay channel provides the cleanest signals at the cost of a lower rate and a reduced acceptance for low p_T J/ψ s, which can be reconstructed only above $p_T \sim 3 \text{ GeV}/c$ in the forward regions and $p_T \sim 7 \text{ GeV}/c$ near midrapidity. The higher mass Υ is unaffected by the muon p_T cut and can thus be measured to zero p_T . In the future, these studies may be extended to the e^+e^- decay channel using the CMS electromagnetic calorimeter. Experience being gained at RHIC in developing triggers for $J/\psi \rightarrow e^+e^-$ and $\Upsilon \rightarrow e^+e^-$ using the STAR electromagnetic calorimeter will be useful for evaluating similar measurements at CMS.

4.3 Jet Reconstruction

Jet reconstruction in heavy ion collisions faces the challenge of identifying localized E_T clusters in the calorimeters on top of a substantial background from the underlying event. In the CMS heavy ion analysis, a modified iterative cone type jet finder is employed that includes an event-by-event subtraction of background energy. In each event, the average tower transverse energy and the dispersion are calculated for each η ring in the barrel and endcap calorimeters and then this average transverse energy plus dispersion are subtracted from all tower energies. Using the corrected tower energies, jets are found with an iterative cone algorithm. Then, the average tower energies and

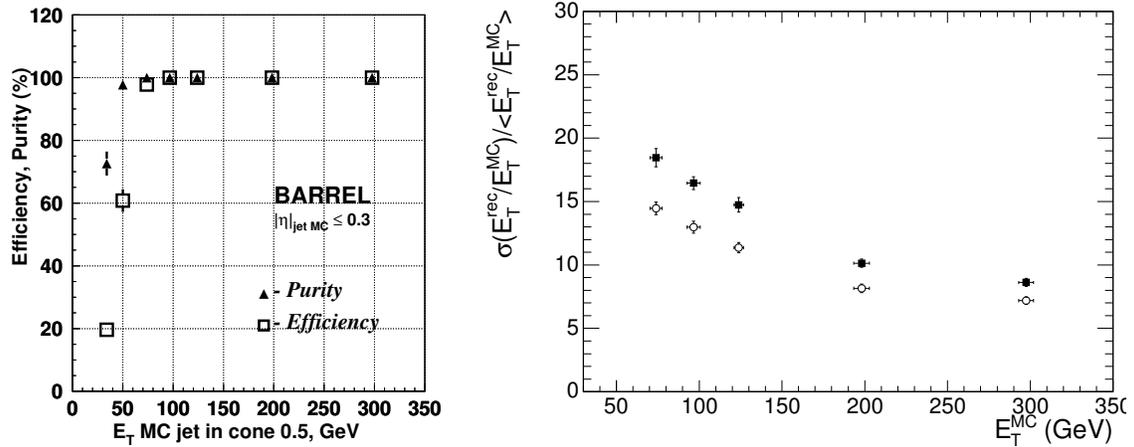


Figure 11: (Left) Jet reconstruction efficiency and purity using barrel calorimeters for PYTHIA generated jets embedded in Pb+Pb events with $dN_{ch}/dy = 5000$. (Right) The resolution of the jet E_T determination in p+p (open symbols) and Pb+Pb (closed symbols), also with $dN_{ch}/dy = 5000$.

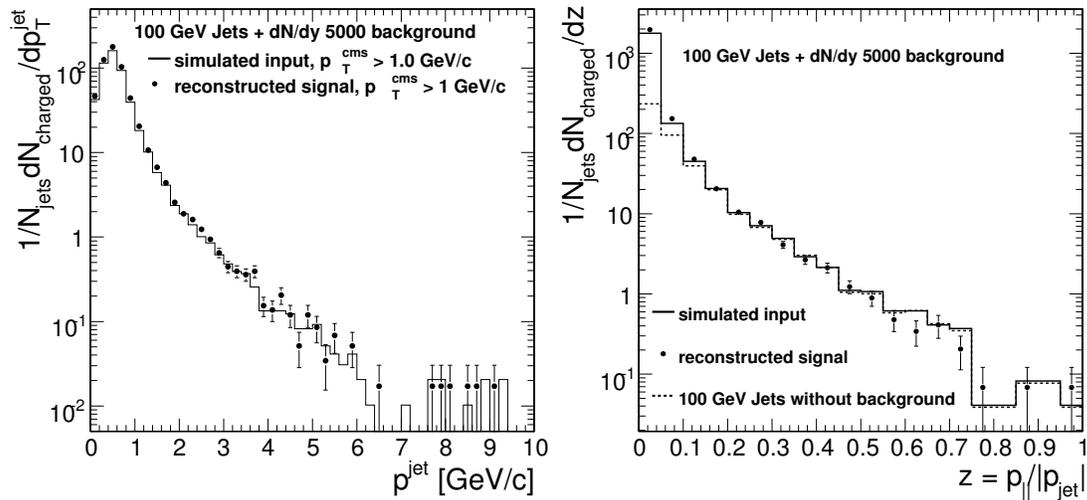


Figure 12: (Left) Input and reconstructed fragmentation function for individual particle momentum transverse to the jet axis for 100 GeV jets embedded in Pb+Pb events with $dN_{ch}/dy = 5000$. (Right) Fragmentation function for longitudinal momentum fraction Under the same conditions. The dashed line shows the distribution for those particles which were actually part of the input jet, rather than the underlying Pb+Pb background.

dispersions are recalculated using only towers outside the jets and the process is repeated with these recorrected tower energies.

This fast method using only calorimeter information is available at the trigger level and already provides excellent reconstruction efficiency and purity as shown in the left panel of Fig. 11. The energy resolution for 100 GeV jets is $\approx 16\%$ (see right panel of Fig. 11) and the jet location resolutions in η and ϕ are 0.028 and 0.032, respectively. The resulting quality of reconstructing the fragmentation function for the momentum of individual particles both transverse and longitudinal to the jet axis is shown in Fig. 12.

Below $E_T \approx 50$ GeV, both the efficiency and purity of full jet reconstruction drop rapidly. To extend our measurements to lower parton energies, we will use the methods developed in the analysis of RHIC data, extracting information on jet-like hadron production on a statistical basis using single-hadron momentum distributions and di-hadron angular correlations. Event characterization using data from the calorimeters will allow us to perform

these measurements as a function of centrality and jet angle relative to the reaction plane. Compared to RHIC, such indirect jet studies at CMS will benefit from larger rapidity acceptance, better statistics, and excellent resolution in both the calorimeters and the momentum from tracking.

4.4 Calorimetry

The most important characteristics of the CMS calorimeters (electromagnetic ECAL and hadronic HCAL), namely the energy resolution and granularity, are listed in Table 5.

Rapidity coverage	$0 < \eta < 1.5$		$1.5 < \eta < 3.0$		$3.0 < \eta < 5.2$
Subdetector	HCAL (HB)	ECAL (EB)	HCAL(HE)	ECAL (EE)	HF
a	1.16	0.027	0.91	0.057	0.77
b	0.05	0.0055	0.05	0.0055	0.05
granularity $\Delta\eta \times \Delta\phi$	0.087 x 0.087	0.0174 x 0.0174	0.087 x 0.087 (except highest η)	changes from 0.0174 x 0.0174 to 0.05x0.05	0.175 x 0.175

Table 5: Energy resolution parameters (using the formula $\sigma/E = b + (a/\sqrt{E})$) and granularity of the CMS calorimeters in the barrel (HB, EB), endcap (HE,EE), and forward (HF) regions. The energy resolution is shown for the total energy of electrons and photons (ECAL) and transverse energy of hadronic jets (HCAL, HF).

The four components of the hadronic calorimeter system, namely the Barrel (HB), Forward (HF), and Endcap (HE) units provide good segmentation and hermeticity, moderate energy resolution, and full azimuthal angle coverage up to $|\eta| = 5$.

The ECAL is built of high density lead tungstate (PbWO_4) crystals, producing a detector which is fast, high resolution, finely segmented, and radiation resistant, all important properties for the LHC environment. Studies have shown the ability of CMS to identify π^0 s in the transverse momentum range 4-20 GeV/c^2 , even for the highest multiplicities expected in Pb+Pb collisions. It is also possible to get statistical separation of neutral particle yields even when the shower-by-shower invariant mass cannot be reconstructed by selecting showers with electromagnetic features. Unless there is very significant jet-quenching, electromagnetic showers will be largely dominated by π^0 s for $p_T < 100 \text{ GeV}/c$.

4.5 Comparison to ALICE

The discussion presented in this section clearly demonstrates the significant capabilities of the CMS detector to perform important measurements for heavy ion physics. Experience at other recent collider facilities (LEP, Fermilab, RHIC) has shown the importance of having multiple detectors addressing a particular physics topic. However, the contribution of CMS to the heavy ion effort will go well beyond a complementary role to ALICE, the dedicated heavy ion detector. Instead, inclusion of CMS provides a number of unique capabilities. These include:

- Acceptance: Significantly more acceptance than ALICE near midrapidity ($|\eta| < 2.5$, full ϕ) for layered detection of charged and neutral hadrons as well as muons, electrons, and photons over a wide range of p_T .
- Resolution: The best mass resolution of any LHC detector for Quarkonia (J/ψ , Υ) leading to clean separation of the various states and improved signal over background. Significantly better charged track momentum resolution over a much wider range of acceptance.
- Calorimetry: Full electromagnetic and hadronic calorimetry for complete jet triggering and reconstruction over a very large solid angle, leading to an increase in the statistical significance of single jet and jet+X channels, where X is another jet, a γ , or a Z^0 .
- Forward coverage: Unparalleled forward physics capabilities including the forward hadronic calorimeter ($3 < |\eta| < 5$), Castor-Totem ($5.5 < |\eta| < 6.6$), and the ZDC ($|\eta| > 8.1$ for neutrals).
- Trigger: The DAQ is capable of delivering almost every Pb+Pb event to the High Level Trigger allowing maximum flexibility to select rare probes. Having the event selection implemented in software makes it easy to update, improve, and modify for a variety of physics programs or to respond to new discoveries.

4.6 Summary

Initial investigations of the CMS detector's properties have found outstanding capabilities for performing important physics analyses in the heavy ion environment. These preliminary studies will be refined and expanded and many additional physics topics will be considered in order to define the full physics potential of the experiment.

5 The CMS DAQ and Trigger System

The demands on the CMS DAQ and trigger system can be summarized as follows: For p+p running at LHC design luminosity, multiple collisions will occur at each bunch crossing, with a frequency of 40 MHz. However, the effective output rate to mass storage is limited to 150 Hz, for technology and cost reasons, corresponding to an output bandwidth of 225 MByte/sec in p+p collisions. Consequently, the trigger system in p+p running has to select less than 10^{-5} of all collision events for permanent storage, while maximizing the sensitivity for new physics. CMS has chosen to meet this challenge with a novel trigger design consisting of only two trigger levels, as shown in Fig. 13. The Level-1 trigger is implemented using custom electronics and inspects events at the full bunch crossing rate. The Level-1 selection reduces the event rate by a factor of 400, to 100 kHz. All further online selection is performed in the High Level Trigger (HLT), using a large cluster of commodity workstations (the “filter farm”). The HLT software environment allows the execution of complex “offline” analysis algorithms, restricted only by the execution time of these algorithms. In CMS p+p running, a reduction of the event rate by a factor of more than 600 in the HLT is required to achieve the design output rate.

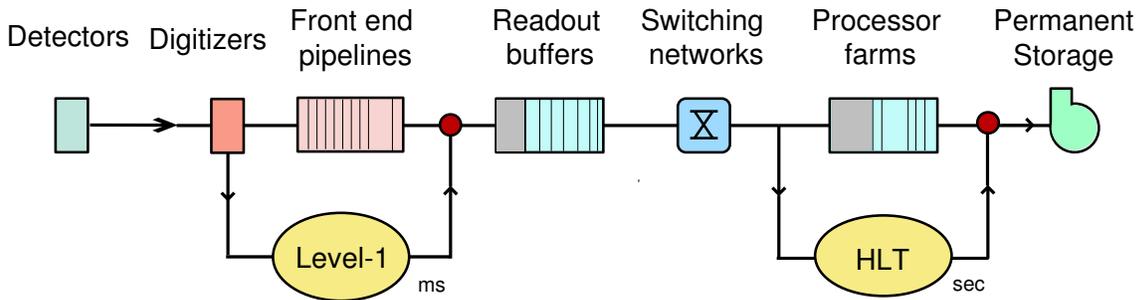


Figure 13: Schematic of the data flow through the CMS DAQ system.

In the following sections, we will show that this trigger architecture can also be used efficiently in Pb+Pb running in CMS. In Pb+Pb mode, the balance of the online rejection will shift from the Level-1 to the HLT trigger. The fast readout of all detector elements, and the high bandwidth up to the HLT, will allow a sequence of HLT trigger algorithms to be run on every heavy ion event, thereby maximizing flexibility and rejection power. Below, we will describe some of the most important physics channels used for event selection. For two of the channels, jets and dimuons, we show the results of detailed timing and performance studies, demonstrating the gain in physics reach due to the HLT. We will concentrate on two Pb+Pb running scenarios: The first one corresponds to Pb+Pb design luminosity with the full DAQ and HLT filter farm in place. The second scenario corresponds to initial LHC running at 1/20th of Pb+Pb design luminosity with 50% of the DAQ and a fraction of the HLT farm available. Following the philosophy of the HLT usage in p+p, we will assume that the present offline algorithms for heavy ion analysis provide the baseline for the online HLT algorithms for event selection.

5.1 Trigger Strategy for Heavy Ion Collisions

In this section, we will discuss the basic trigger strategy foreseen in heavy ion running of CMS. To understand the choice of this strategy, we need to introduce some of the basic constraints. At Pb+Pb design luminosity, the initial true collision rate at the beginning of a store is expected to be close to 8 kHz. Assuming collisions at three detectors, the collision rate will drop sharply over the time of a store, to an average of about 3 kHz. The maximal rate of 8 kHz for Pb+Pb is much smaller than the 100 kHz HLT input rate for p+p running. On average, minimum bias Pb+Pb events are larger than p+p events by only a small factor, because a typical p+p “event” at design luminosity actually consists of ≈ 25 superimposed p+p collisions. Therefore, the HLT input bandwidth, as determined by the p+p requirements, is sufficient to send *all* Pb+Pb events to the HLT, even at maximum collision rate. However, it will be necessary to tune the L1 trigger to reject backgrounds such as beam-gas interactions (see Section 7).

Experience from studies of heavy ion collisions at the SPS and RHIC shows that there is no “simple” criterion for rejecting events based on global characteristics like multiplicity or average transverse energy. Rather, it was found that for essentially all observables, studies as a function of collision centrality are critical for extracting the

underlying physics. Therefore, the basic trigger strategy for heavy ion collisions has to be the efficient identification of any potentially interesting signature for any given input event. Experience also shows that the rejection of background events, such as those caused by beam-gas collisions, early on in the triggering chain is crucial for providing a sufficiently clean, low rate environment for the higher level trigger. The main purpose of the Level-1 trigger in CMS heavy ion running will be to provide such a clean discrimination of true heavy ion collisions and to provide seed objects such as high p_T muon candidates as input to the HLT algorithms. No significant rejection of Pb+Pb collisions at the Level-1 trigger is foreseen.

The possible gain in physics reach by the HLT, in comparison to simply collecting minimum bias events, is determined by the ratio of the collision rate to the rate of events written to mass storage, which is 225 MByte/sec. This translates into an event rate of 10-100 Hz, based on estimates of the heavy ion event size that will be discussed below. This output rate is not only limited by the available mass storage technology, but also by the limits on the available offline analysis resources. It can be argued that it is more efficient to invest resources into a high quality online trigger scheme, rather than into the offline handling and storage of a poorly selected data stream.

The rates quoted above suggest a maximal average gain due to triggering of a factor of 30 to 300 in the statistics of rare probes. In practical terms, this gain in effective rate to tape will make the difference between measuring the overall yield of a rare process compared to studying the same process as a function of impact parameter and reaction plane or the difference between making a measurement in one heavy ion run at the LHC versus making the same measurement over a span of 20 years. Simply put, triggering at the LHC will be crucial. The flexibility of the HLT system will allow us to allocate bandwidth not just to certain trigger channels, but differentially as a function of rapidity and p_T of the trigger object and as a function of collision centrality, thereby maximizing the overall physics reach of our measurements.

The basic trigger strategy in Pb+Pb running can be summarized as follows: Every Pb+Pb collision in our interaction region identified by the Level-1 trigger will be sent to the HLT filter farm. At the HLT, the full event information will be available for each event. All rejection of Pb+Pb collisions will be based on the outcome of HLT trigger algorithms that are either identical with the corresponding offline algorithms or optimized versions of the offline algorithms. This strategy relies on the fact that the HLT in its final configuration will provide sufficient input bandwidth to accept all collision events, even at the maximum Pb+Pb collision rate. The strategy also requires that the algorithms can be executed quickly enough. For comparison with the expected HLT CPU budget per event, we will show the expected gain in physics reach for selected algorithms, and present timing measurements of those algorithms. Overall, the selectivity that can be achieved by the trigger depends on the availability of sufficient CPU resources to execute the algorithms, possibly triggered by a Level-1 seed, and on the efficiency, acceptance, and background level of each particular algorithm. The output event rate is determined by the ratio of the average size of the selected events to the “allowed” bandwidth to tape. Detailed discussion of these issues is presented in Section 8.

5.2 Trigger Channels in Heavy Ion Collisions

The results from RHIC have clearly pointed out the importance of probes at intermediate and high p_T , such as leading hadrons, di-hadron correlations, and spectra and azimuthal distributions of hadrons carrying charm. At the LHC, studies of high p_T hadron production can be extended from the present RHIC statistical limit of $p_T < 20$ GeV/c to transverse energies of several 100 GeV. Similarly, studies of open and hidden heavy-flavor physics at the LHC will be extended to include b -quark production. Studies of vector bosons and fully formed jets in heavy ion collisions will become possible for the first time. Like the current studies at RHIC, the corresponding studies at the LHC will ultimately be limited by the statistics given by production and rates and integrated luminosity. The Level-1 and HLT triggers are essential for maximizing the physics reach of CMS within these constraints.

As shown in previous sections, the high granularity of the CMS inner silicon tracker allows the reconstruction of a large fraction of the produced hadrons even in central Pb+Pb collisions. The current execution time of the track reconstruction algorithm (~ 1200 sec for a central event) prohibits running the full reconstruction on each event at the HLT level. While future studies will explore the possibility of running regional tracking to detect high p_T hadrons, including those from heavy flavor decays, the present trigger studies focus on channels related to calorimeter and muon chamber triggers. Figure 14 summarizes the production cross sections of some of the relevant physics channels, found using version 6.326 of the Pythia event generator [48]. The range of cross sections

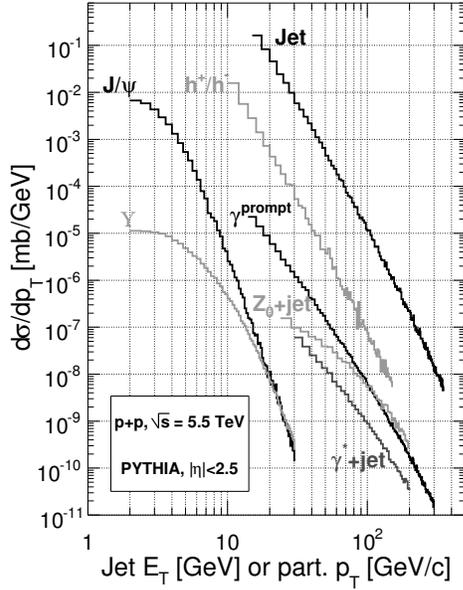


Figure 14: Production cross sections of some of the relevant physics channels in p+p collisions at $\sqrt{s_{NN}} = 5.5$ TeV, as calculated with the PYTHIA event generator.

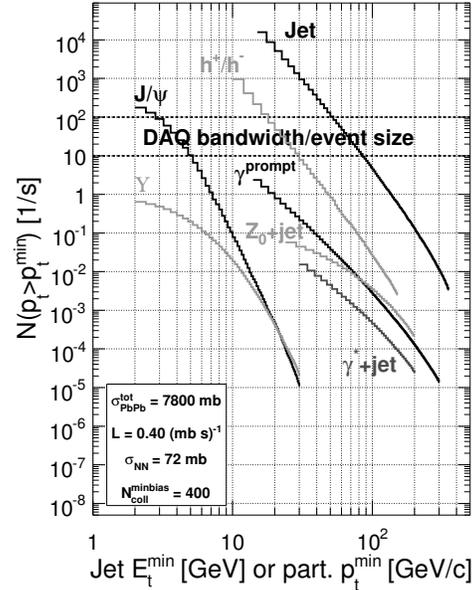


Figure 15: Production rates in minimum bias Pb+Pb collisions at $\sqrt{s_{NN}} = 5.5$ TeV corresponding to the cross sections from Fig. 14 scaled by N_{coll} and design luminosity.

Signal	Threshold	Yield
h^+/h^-	$p_T > 25$ GeV/c	2.1×10^7
h^+/h^-	$p_T > 50$ GeV/c	8.2×10^5
jet	$E_T > 100$ GeV	4.4×10^6
jet	$E_T > 200$ GeV	1.3×10^5
γ^{prompt}	$p_T > 25$ GeV/c	4.3×10^5
J/ψ	–	3.7×10^8
Υ	–	8.7×10^5
γ^*	–	1.5×10^4
Z^0	–	1.3×10^5

Table 6: Integrated yield of selected channels assuming design luminosity and a heavy ion run of 10^6 sec. For the quarkonia and boson channels, the branching ratio into dimuons has not been applied.

of interest extends over 10 orders of magnitude, leading to the corresponding variation in production rates shown in Fig. 15. For this figure, the production rates were calculated for an average luminosity of $L = 4 \times 10^{26} \text{ cm}^{-2} \text{ s}^{-1}$ and minimum bias collisions which have an average number of binary nucleon-nucleon collisions of $N_{coll} = 400$. The corresponding Pb+Pb collision rate is 3 kHz. Also indicated on Fig. 15 is the expected limit on event rates to tape (see Section 8.5), up to a factor of 300 smaller than the average Pb+Pb collision rate. Table 6 shows the integrated production yields for the same channels for one nominal CMS heavy ion run, assuming design luminosity, a total up-time of 10^6 sec, and an ideal detector with infinite bandwidth. Clearly, differential studies of Υ production and jets at high E_T , as well as measurements of γ -jet correlations, will require a highly efficient trigger that selects a large fraction of the corresponding interesting events for tape storage. This is particularly true for studies of such observables as a function of centrality and event-plane. The purpose of the heavy ion trigger will be the allocation of the available output bandwidth to a selection of trigger channels, such that the overall physics impact of the CMS heavy ion program is maximized.

To illustrate the importance of the trigger, we have studied in detail its performance for dimuon and jet measurements. Studies for other trigger channels will be forthcoming. The results presented in Section 8 clearly show the

power of the CMS filter farm when applied to heavy ion collisions, but also demonstrate the need for substantial HLT computing resources in heavy ion running. In the following sections, we will describe the architecture of the CMS DAQ system and the HLT filter farm in detail, and discuss applications of the Level-1 trigger system in heavy ion collisions. Finally, we will present simulations of the HLT performance in heavy ion collisions, including studies of the timing performance and event size distributions from full GEANT-4 based simulations.

6 Overview of the Data Acquisition System

In CMS, the online inspection of events is divided into two steps, as illustrated in Fig. 13 in Section 5. The Level-1 trigger uses local data from the calorimeter and muon systems to make electron/photon, jet and energy sum, and muon triggers. The decision is sent to the front-end detector electronics after a latency of $3 \mu\text{s}$. Events fulfilling the requirements are then transferred to the HLT, where further selection takes place. The rejection rate at Level-1 needs to be sufficiently large to reduce the accepted event rate to about 100 kHz for p+p to match the High Level Trigger (HLT) input bandwidth.

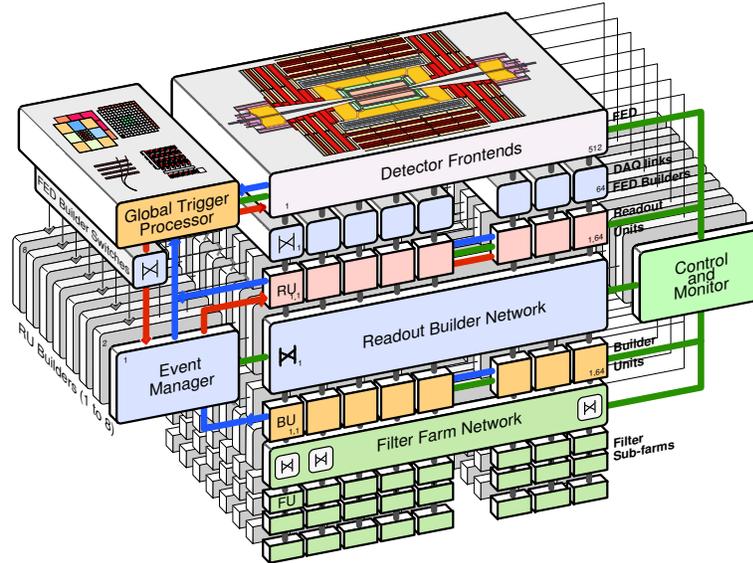


Figure 16: Architecture of the DAQ system. Note the eight slices (arrayed from back to front in the illustration) each consisting of an event building network followed by filter farm processors.

6.1 Data Acquisition System Architecture

The components of the readout system are shown schematically in Fig. 16. In the initial stage (labeled “Detector Frontends” in the figure), data are formatted in ~ 650 Front-end Driver modules which are then read out in parallel by 512 Front-end Readout Links (FRL). The data from eight FRLs are combined into larger event fragments which are sent on to the Readout Unit (RU) buffers. The RU buffers and everything that follows them are organized hierarchically into eight slices as shown. The 64 merged FRL fragments from a given event all go to the matching number of RUs in one of the slices and are subsequently combined and stored as complete events. The final section of each slice, the “Filter Farm”, is further compartmentalized into 8 subfarms, each mounted in a single rack and consisting of 8 Builder Unit (BU) processors and 24 Filter Units (FU), in which the HLT processing is performed. Therefore, a full DAQ slice would contain 64 RUs and BUs, and 192 FUs.

The advantage of this hierarchical and subdivided organization is twofold. On the one hand, it allows a staged deployment of the total computational power required. On the other hand, the segmentation simplifies the management, as well as the dispatching of control, monitor, and error messages. It also provides a natural system partitioning which matches the architecture of the rest of the DAQ. As for the output data produced by the farm, this grouping improves the robustness of the system making it possible to keep the output streams separated at the level of the subfarm, thus decreasing the risk of data loss (and experiment dead-time) due to unavailability or fault in the central data storage services. In case of failure of a subfarm causing data corruption, the amount of lost data is also minimized.

Each FRL is designed to sustain a 200 MBytes/s continuous rate, with peaks up to 400 MBytes/s. For p+p running, the event fragment size read by each FRL is ~ 3 kBytes and so the total event size is expected to be 1.5 Mbyte. For heavy ion running, the data sizes are dominated by the tracker, which is expected to produce on average 1.5 Mbytes of data with central events reaching 8.5 Mbytes or more depending on the event centrality and the level of zero

suppression. The full event sizes are expected to be 2.5 Mbytes and 9.5 MBytes for average and central Pb+Pb, respectively. Readout of the tracker is performed by 250 FRL modules so that typical fragment sizes of the tracker data are expected to be of the order of 10 kBytes, well within the limits of module memory and the switching system.

6.2 Filter Unit

Each FU node is an individual PC with limited local resources (both memory and disk space) but with access to external, shared resources. It is part of a distributed computing environment consisting of the Filter Farm, the Run Control and Monitoring System, and the Computing Services. Shared resources include boot servers, configuration and file servers, and databases containing software, configurations, and other information. The FUs provide the CPU power required to select events and, in turn, determine the physics selectivity of the DAQ chain. Each FU makes indirect transactions to the Computing Services and Run Control System through the Subfarm Manager (SM) which connects to the central computing and storage facilities in order to handle the distribution of code, configuration, and control commands as well as collecting monitor data.

The subfarms need to have local storage capacity to hold the data corresponding to about one machine fill, on the order of 100 GBytes. In addition, it must be possible for the Filter Farm to continue working while the connection to the central computing services is unavailable. Therefore, local mass storage is provided by each individual node of the subfarm to store accepted events. Subsequent transfers to the central computing services can be conveniently sized and scheduled, for example on a run-by-run or fill-by-fill basis.

6.3 The HLT Application

The High Level Trigger selection algorithms run as an application which consists of an event loop, a series of reconstruction units performing parts of the reconstruction necessary to reach the trigger decision, and a number of filter modules that apply the selection criteria. According to a policy set out in the configuration of the HLT, certain modules can break the processing and unconditionally reject the event (HLT veto). In all other cases, an event is processed until the full chain of filter modules has reached a decision and is accepted if it passes any one of the filter modules. After an accept or reject, the filter application may perform additional calculations before dispatching a response to the framework and then the HLT application is ready for a new event. Development of the specific reconstruction and selection algorithms for the HLT application is one of the primary responsibilities of the heavy ion physicists working in CMS.

6.4 DAQ Deployment Schedule

The hierarchical DAQ architecture allows staged deployment of the full system. The schedule is designed to match the lower event building bandwidth required at the start of the program. At LHC machine startup in the middle of 2007, we expect to have the complete first stage, the FED builder, and up to 50% of the infrastructure of the second stage, the RU builder, fully operational. This will include PC processors for the RUs and the BUs. The initial number of Filter Farm PCs available is likely to be small, on the order of 100 out of the full complement of 1536. As the luminosity of the p+p beams increases and the complex Pb+Pb events need to be processed, the experiment will be adding Filter Units. The 4 initial slices, corresponding to 50% deployment of the DAQ with a total of 768 FUs including the first shipment of PCs requested in this proposal, are expected to be available by the end of 2008. The full DAQ, including the remaining CPUs discussed in this proposal, is scheduled to be deployed by the end of 2009.

7 Level-1 Trigger

7.1 Introduction

The design of the Level-1 trigger was optimized for p+p data taking. Specific emphasis is placed on triggering on relatively rare events such as high-transverse energy jets, missing transverse energy, and muons. The main objective of Level-1 is to reduce the flow of events from the 1 GHz collision rate, or 40 MHz beam crossing rate, ($p+p$ at $L = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$) down to 100 kHz, which is an acceptable rate to the High-Level Trigger (HLT).

The Level-1 trigger is organized into two main subsystems: the calorimeter and muon triggers which are combined to generate the global trigger, as illustrated in Fig. 17. The calorimeter trigger begins with the trigger tower energy sums in the EM and hadronic calorimeters. Towers with an energy above a preset threshold are used by the regional calorimeter trigger to find candidate jets, tau leptons, electrons, and photons. The candidates are then sorted in energy, with only the top four of each type retained for final processing in the global Level-1 trigger. The muon trigger utilizes the three muon subsystems in CMS: drift tubes (DT), cathode strip chambers (CSC), and resistive plate chambers (RPC). Muon candidates are formed from either the DT/CSC system, the RPC, or a combination of both. Again the top four candidates (in p_T and η, ϕ quality) are retained for the global trigger decision. In order to achieve an acceptable data output rate, the global trigger applies energy cuts and prescaling factors to accept only data relevant to different specific physics scenarios, such as high energy jets, dijets, or jet+electron. Events accepted at this time are allowed to proceed to the HLT for the write-to-tape decision.

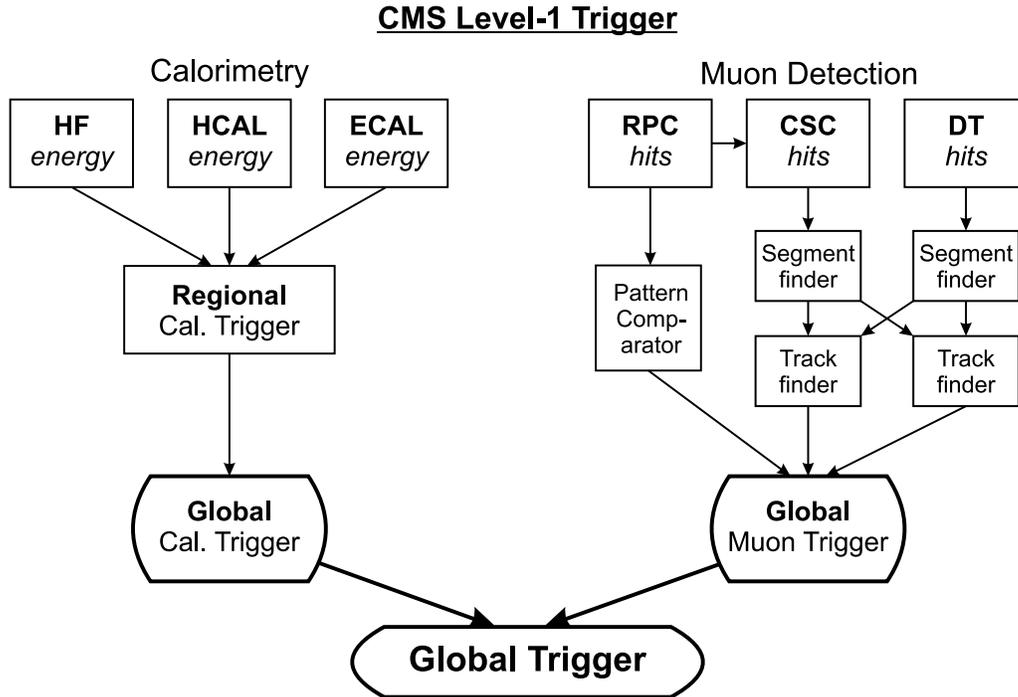


Figure 17: Simplified flow chart showing the major components of the CMS Level-1 trigger system.

For p+p collisions, a subset of minimum bias data will also be taken for reference and specific physics measurements. In the heavy ion program, the Level-1 triggering requirements for physics output will not be too different from the minimum bias p+p setup, although the details of the implementation in Level-1 will likely require some modification. For example, additional triggering based on the ZDCs will be possible through use of trigger bits providing time-difference and detector-hit information.

7.2 Minimum Bias Triggers

The baseline results for both p+p and Pb+Pb collisions will be derived from minimum bias data. A majority of this data will likely be taken during the early, low luminosity running of the LHC, both to minimize pileup concerns

(primarily in the p+p case) and for practical reasons of detector commissioning and calibrations. Centrality dependence is an important component of many heavy ion analyses and, therefore, the minimum bias data is critical. Attaining a Level-1 minimum bias efficiency as close as possible to 100% is highly desirable for a confident understanding of the centrality and, thereby, minimization of the uncertainty in the event classification.

Peripheral Pb+Pb collisions can be approximated by p+p events. Thus, development of high efficiency triggering for p+p should lead to a high quality minimum bias trigger for heavy ions. Work is in progress to study triggering on p+p minimum bias collisions with an expected efficiency at Level-1 on the order of 80-90%. Initial studies have supported a RHIC-style minimum bias trigger, in which only the forward regions of the hadron calorimeters will be used. A simple counting of the hit calorimeter towers in these regions could form the basis of the trigger and rudimentary beam-gas background rejection (see Fig. 18). Note that the beam-gas background is expected to be a smaller effect than seen at RHIC due to the improved vacuum in the LHC. Work is continuing to optimize this event selection for Pb+Pb.

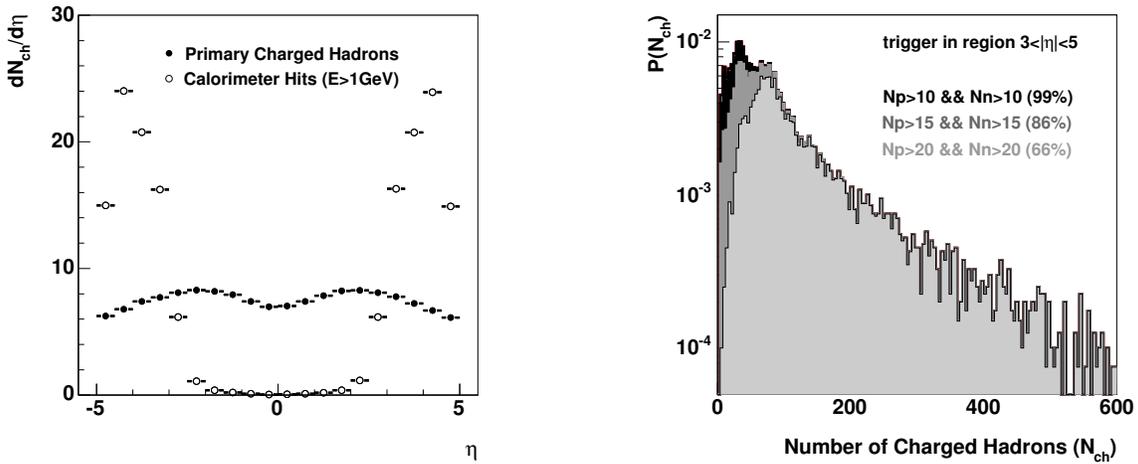


Figure 18: (Left) The number of charged particles in 14 TeV p+p collisions (closed symbols) along with the simulated number of detected particles in the hadronic calorimeter (open symbols). The difference in the forward region is predominantly from contributions of neutral particles and secondary particles produced from interactions with the detector material. (Right) Simulated efficiencies when triggering with the forward hadronic calorimeters, where the labels N_p and N_n refer to valid hits above threshold on the “positive η ” and “negative η ” side of the collision point.

7.3 Jets at Level-1

To enhance the recorded data with rare events, Level-1 can be used to select events with high transverse energy in a narrow angular region - one of the typical signatures for a jet. As jets will allow the probing of the early stages of the collision, triggering on events with such characteristics is desirable. Algorithms already exist to trigger on jets at Level-1 for p+p collisions. However, jets that clearly stand out in the lower multiplicity environment of p+p will lie on top of a baseline from the underlying event in Pb+Pb collisions. As an example, the turn-on for 100% finding efficiency of jets at CMS Level-1 in p+p collisions has been estimated to be $E_T \sim 30$ GeV. However, for Pb+Pb collisions, a conservative estimate for this threshold is $E_T \sim 80$ GeV (see, for example, Fig. 11). In this case, further improvement in efficiency and purity will be possible at the HLT level, and most likely the Level-1 trigger system would not actively accept or reject such events but instead simply pass on $\eta - \phi$ “jet seeds” to the HLT for further analysis.

7.4 Muons at Level-1

Additional analyses such as detection of muons from J/ψ and Υ -family decays can profit from the event preselection capabilities of Level-1. Muon triggering is also essential for analyses requiring jet-flavor tagging, providing a means for differential study of partonic energy loss mechanisms. Most of the information used to find muons at Level-1 is derived from the DT, CSC, and RPC detectors, used either singly or in combination. Position information is read-out for each detector, yielding a rough transverse momentum measurement.

The performance of the Level-1 Muon trigger is affected by the detectors situated in front of the muon chambers. The absorption of the muons in the calorimeter material limit the p_T range to above 3.5–4.0 GeV for 90% efficiency in the central barrel ($|\eta| < 1.0$) (3.0 GeV for 50% efficiency). The lower bound is reduced to 1.5 GeV in the more forward regions ($|\eta| \sim 2$). Single muon triggering at Level-1 has also been found to enhance the acceptance for quarkonia measurements in the heavy ion environment[49].

7.5 Level-1 Outlook

In summary, the Level-1 trigger is an important tool for enhancing the physics reach of CMS. Work is in progress to understand the details of the critical minimum bias triggering for both p+p and Pb+Pb collisions, with the goal of keeping as many real events as possible while not overloading the HLT with background events such as beam-gas. Analysis of certain rare processes may benefit from additional Level-1 triggers whose ability to perform satisfactorily during heavy ion runs will be investigated. However, the bulk of improvement in triggering in the heavy ion environment will come from the HLT and its capability for more sophisticated algorithms.

8 High Level Trigger Performance

As outlined in Section 5, the trigger strategy in Pb+Pb foresees running reconstruction algorithms on all Pb+Pb collisions in the HLT filter farm, using the full event information. In this section, we present studies of the HLT performance in the dimuon and jet channels. The physics performance for a minimum bias trigger will be compared to that of a trigger system allocating fractional bandwidth to various trigger channels implemented using the Level-1 and HLT. For these studies we used three sets of fully simulated Pb+Pb events corresponding to $b = 0, 9,$ and 12 fm impact parameter, processed with the CMS Geant-4 based ORCA simulation package. Timing measurements used the standard heavy ion offline algorithms in the ORCA software package run on standard PCs with 1.8 GHz Opteron CPUs and 4 GByte RAM.

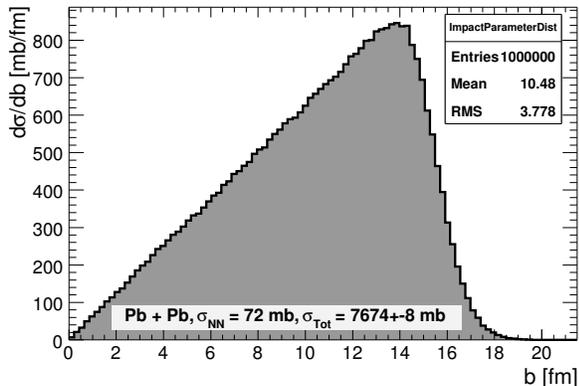


Figure 19: Impact parameter distribution calculated with a Monte Carlo Glauber simulation.

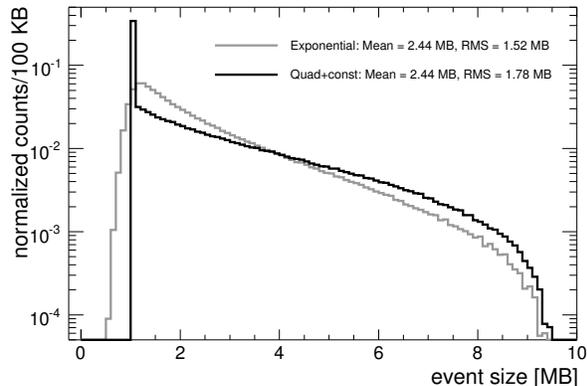


Figure 20: Minimum bias event size distributions obtained by folding two parameterizations of the event size distribution with the impact parameter distribution.

8.1 Event Size

The possible gain in statistics for rare probes provided by the online trigger system is determined by the ratio of collision rate to rate to tape which is 225 MByte/sec, thus requiring an estimate of the average event size to obtain a rate of events. The event size was obtained by writing simulated raw data, including hits from background and secondary particles, in the form of CMS standard “Digis”. The average mid-rapidity charged hadron density for these three data sets is $dN_{ch}/d\eta|_{\eta=0} = 3300, 575,$ and $65,$ respectively. For all hits, online pedestal subtraction and zero-suppression are performed, but no further suppression or encoding of the hit information was attempted. The event-size was found to be approximately linear in the overall multiplicity, with an offset that was adjusted to match the standard p+p event-size of 1.5 MByte at design luminosity for equal multiplicity. The resulting event size vs collision centrality distribution was fit with two parameterizations as a function of impact parameter, $b,$ and then folded with the impact parameter distribution obtained from a Glauber calculation (see Fig. 19). The resulting event-size distributions for minimum-bias Pb+Pb running are shown in Fig. 20, leading to the estimated average event size of ≈ 2.5 MByte per minimum bias event and ranging up to ≈ 9.5 MByte for the most central events. With an average rate to tape of 225 MByte/sec, the corresponding average minimum bias event rate to tape is 90 Hz. Clearly, this estimate depends on the actual multiplicity of produced hadrons in Pb+Pb collisions at the LHC, for which the range of predictions and extrapolations suggest an uncertainty of about a factor of three. Experience suggests that the background levels are likely to be somewhat underestimated by simulations, which is partially taken into account in the p+p estimate to which these studies were matched. In later running periods, this will likely be more than offset by more sophisticated encoding of the hit information. In the PHOBOS silicon detectors, lossless Huffman data compression yielded a factor of four reduction in the event size for the silicon subsystem, but required a significant amount of online computing resources and a good online calibration of all detector channels. The possible use of the HLT system for performing such a compression in CMS will be studied in the future.

For the studies presented below, we will use the event size distribution as shown in Fig. 20. Including all of the uncertainties, we expect a rate of events to mass storage between 10 and 100 Hz. A large part of this uncertainty

will only be resolved once the first data at LHC are taken, which underscores the need for a flexible high level trigger scheme.

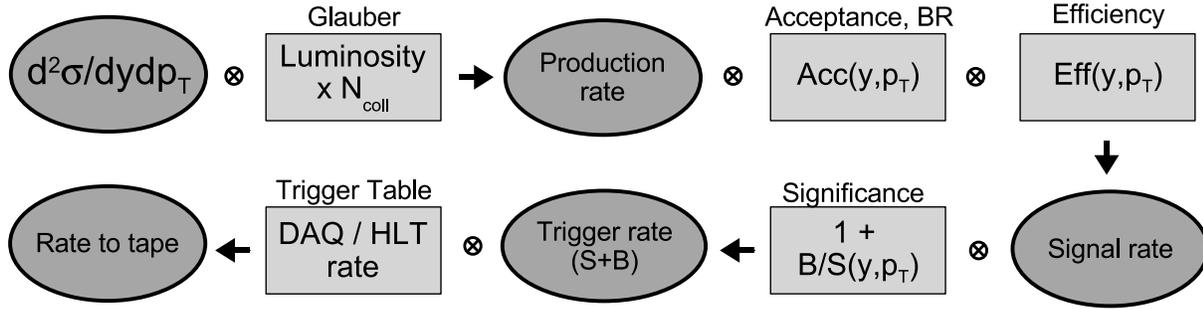


Figure 21: Schematic description of the HLT simulation chain. Input cross sections are obtained from Pythia v.6.326. Acceptance, efficiency, and background levels are parametrized from full Geant-4 based simulations. See text for detailed descriptions of each of the steps.

8.2 Trigger Timing Studies

The capabilities of the HLT are largely determined by the time required to process each event and are therefore strongly dependent on the CPU resources available in the filter farm. As discussed in Section 6, the online farm consists of about 1500 compute servers. Because of the expected increase in CPU performance/cost over the coming years, the detailed specifications of these servers is not precisely known. For the purpose of our present studies, we will assume that the servers purchased for 2008 running will have dual CPUs with quad cores, with a performance per core comparable to 1.5 times the performance of the 1.8 GHz Opteron CPUs used now. This estimate is conservative since the actual trigger PCs will be purchased no earlier than a year and a half from now. Timing measurements will be quoted in units of CPU seconds for the present CPUs, for which the time budget per event for the full HLT filter farm will be ~ 1.5 sec at the beginning of each store, and ~ 4 sec averaged over the duration of the store.

Like the event size studies, the timing studies were performed using the full ORCA based simulations and the present offline algorithms for jet finding and dimuon reconstruction. As of now, these algorithms have been optimized for reconstruction efficiency and background rejection, but not for timing performance. Based on generator level studies with slightly simplified geometries, and comparisons to the optimized p+p algorithms, we expect that significant gains relative to the timing performance shown below can still be achieved. However, these gains will be partially offset by the CPU resources needed for running the algorithms for additional trigger channels as shown in Table 7. These additional algorithms have so far only been tested in generator level studies which do not allow detailed timing measurements.

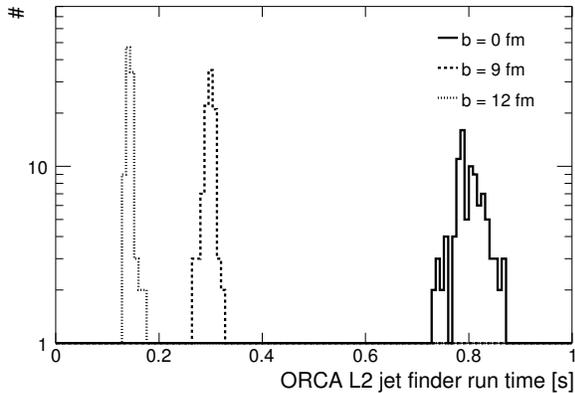


Figure 22: Execution time distribution of the $L2$ jet finding algorithm for $b = 0, 9,$ and 12 fm.

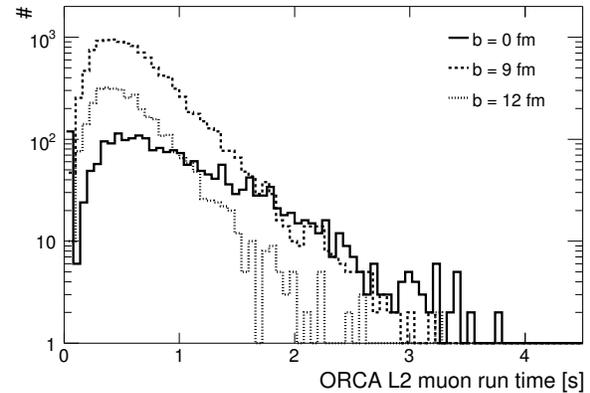


Figure 23: Execution time distribution of the $L2$ muon finding algorithm for $b = 0, 9,$ and 12 fm.

8.2.1 Jet Finder Timing

The online jet finding algorithm, like the present offline algorithm studied here, will be run on so-called trigger towers. Each trigger tower combines one hadron calorimeter cell with about 25 electromagnetic calorimeter cells, for a total of 4176 trigger towers. In Fig. 22, we show the execution time distribution of the jet finder for the three different impact parameter bins. For all bins, the distribution time is nearly Gaussian, with averages of 820, 320, and 164 msec for $b = 0, 9,$ and 12 fm, respectively. Averaging over the Glauber impact parameter distribution, the estimated average execution time is $\langle t \rangle = 250 \pm 25$ msec. The error is dominated by the uncertainty in the functional form of the impact parameter dependence of the execution time, due to the small number of points. Therefore, if it is executed for every event, the jet finding algorithm is expected to use between 5 and 15% of the CPU budget.

8.2.2 Muon Finder Timing

The muon finder consists of three different algorithms. The first part, $L1$, is executed in Level-1 for each event, producing a list of muon candidates. This algorithm is closely based on the corresponding algorithm for p+p although the cuts have been adjusted to allow for a larger acceptance at low p_T . The second part of the algorithm, $L2$, is also executed on all events, this time in the HLT. The execution time distribution for this algorithm is shown in Fig. 23 for the three impact parameter bins, giving execution times with averages of 706, 95, and 10 msec for $b = 0, 9,$ and 12 fm, respectively. Averaging over impact parameter yields an estimated average execution time of $\langle t \rangle = 80 \pm 10$ msec. Correspondingly, the $L2$ execution uses 1 to 5% of the HLT CPU budget.

The third part of the algorithm, $L3$, is run in the HLT on events with at least two muon candidates found by either $L1$ or $L2$. Our measurements show that, averaged over impact parameter, $L3$ will only be called for $2 \pm 1\%$ of all events. However, as $L3$ requires tracking in the silicon detector, the execution time is significantly longer than for any other algorithm studied here and shows a very steeply increasing dependence on multiplicity. The $L3$ execution time is also found to be linear in the number of muons candidates from $L2$. Averaging over the impact parameter distribution of inputs selected by $L1$ or $L2$, we find an execution time at the $L3$ level of 1750 ± 200 seconds. Per minimum bias event, this is an order of magnitude larger than the available CPU budget. As mentioned before, the timing performance of the $L3$ algorithm has not been optimized. We expect, based on the steepness of the multiplicity dependence, that significant gains in performance are possible with further development, while largely preserving the additional rejection factor of 4 provided by $L3$. Still, at design luminosity, the HLT muon finding algorithms are likely to consume a significant fraction of the time budget even for the full configuration of the filter farm.

8.3 Trigger Signal Rates

To evaluate the physics reach of the CMS detector, the production rates shown in Fig. 15 need to be translated into rates of observed events in the detector and finally into the rates to mass storage for the physics processes of interest. A straightforward measurement of the rates seen in full simulations is presently not feasible. Available CPU resources limit the sizes of typical MC data sets to about 100k, compared to a total integrated number of collisions of several times 10^9 for a nominal LHC heavy ion run. As we are only interested in triggering on rare probes, we can instead use the approach shown schematically in Fig. 21 and described in the following sections. The production rates shown in Fig. 15 were increased from p+p to Pb+Pb assuming scaling with the number of nucleon-nucleon collisions, and then multiplied with branching ratios, where appropriate, to obtain the rate of signal events in the form of three-dimensional histograms in $y, p_T,$ and impact parameter.

The acceptance and efficiency of the full offline algorithms for simulated events were parametrized in corresponding (y, p_T) histograms. Multiplication of the production rate histograms with acceptance and efficiency histograms yields the *trigger signal rate* shown in Fig. 24. Note that the rates lie below the output limit for J/ψ and Υ production, while the rate of possible jet triggers exceeds the limit for trigger thresholds of $E_T < 40$ GeV.

8.4 Trigger Background Rates

Offline studies show that the jet and dimuon triggers will include a substantial fraction of background events, depending on the chosen E_T and p_T threshold. In addition, the dimuon analysis requires a certain number of

events outside of the J/ψ and Υ mass windows to allow a reliable estimate of the coincidental background at the quarkonium masses. To include these backgrounds in our rate estimate, we parametrized the fraction of background to signal events, B/S , from the full offline simulations as a function of p_T (E_T) and multiplied the trigger signal rates discussed previously with a factor $(1 + B/S)$, to obtain the actual trigger rates which are shown in Fig. 25.

8.5 Output Rates

As Fig. 25 shows, the possible trigger rates from the dimuon and jet channels alone far exceed the limit of 10–100 Hz, if no p_T threshold for the trigger is applied. Additional bandwidth will be needed for composite channels such as γ +jet, for a trigger on ultra-peripheral events in studies of forward physics, and for minimum bias events. The available output bandwidth therefore needs to be allocated to the various trigger channels by applying trigger thresholds and scale-downs, such that the overall physics reach of the experiment is maximized. In allocating the bandwidth, it is important to account for the fact that the average size for events satisfying a given trigger condition will typically be significantly larger than that for minimum bias events. With the obvious exception of ultraperipheral triggers, the assumption that rates of the interesting processes scale with the number of nucleon-nucleon collisions heavily biases the accepted distributions towards more central collisions. In Table 7 we show a strawman trigger table which illustrates the possible allocation of the total bandwidth to individual trigger channels. Clearly, this table will have to be optimized to maximize the scientific output of the CMS heavy ion program. The resulting rates to tape can be seen in Fig. 26, which shows rates of events to tape for minimum bias running (no event selection in HLT) in comparison to those for event selection using the high level trigger with the bandwidth allocation given in Table 7.

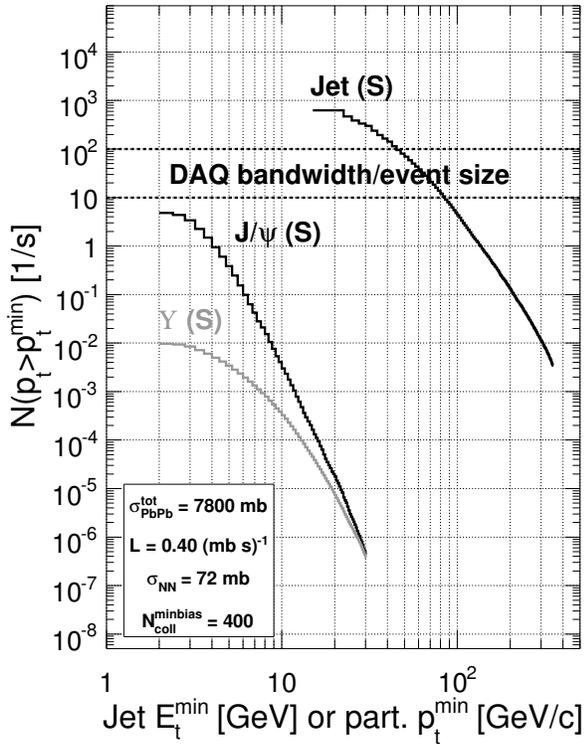


Figure 24: Average trigger signal rates for J/ψ , Υ , and jet production at design luminosity. The horizontal band shows the 10 to 100 Hz estimated DAQ event rate limit.

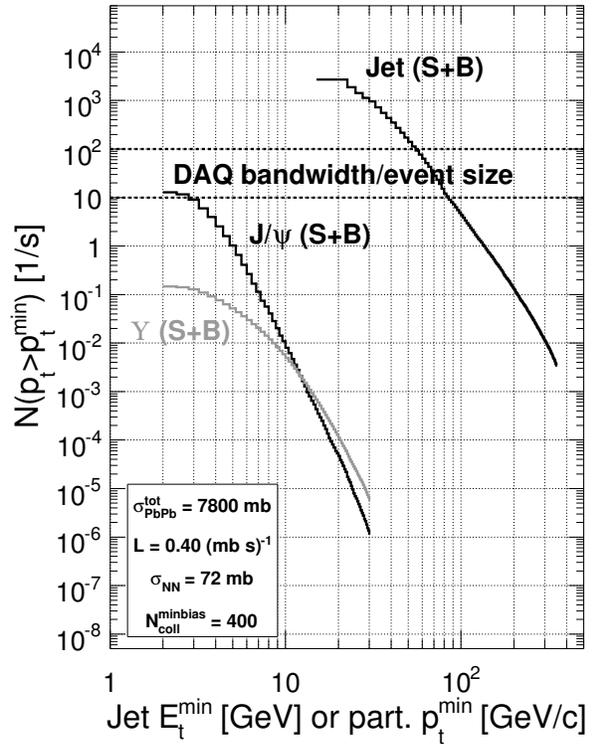


Figure 25: Average trigger rates including both signal and background for J/ψ , Υ , and jet production at design luminosity.

Channel	Threshold	Scale-down	Bandwidth [MByte/sec]	Event size [MB]
Min.bias	—	1	33.75 (15%)	2.5
Jet	100 GeV	1	24.75 (11%)	5.8
Jet	75 GeV	3	27 (12%)	5.7
Jet	50 GeV	25	27 (12%)	5.4
J/ψ	0 GeV/c	1	67.5 (30%)	4.9
Υ	0 GeV/c	1	2.25 (1%)	4.9
γ^{prompt}	10 GeV	1	40.5 (18%)	5.8
UPC/forward coll.	—	1	2.25 (1%)	1

Table 7: Strawman trigger table for running at design luminosity, assigning fractions of the total bandwidth (225 MByte/sec) to individual trigger channels. The last column shows the average event size for each of the trigger streams.

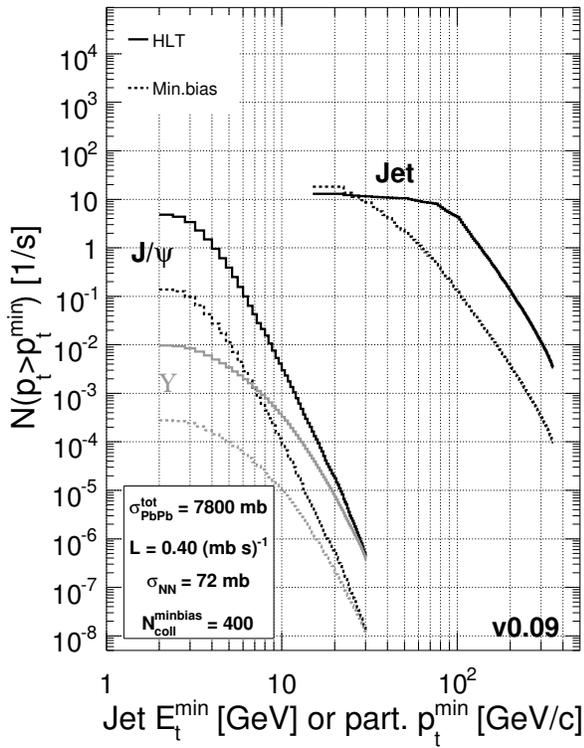


Figure 26: HLT J/ψ , Υ , and jet rates recorded to tape using the trigger settings in Table 7 compared to those for a minimum bias trigger.

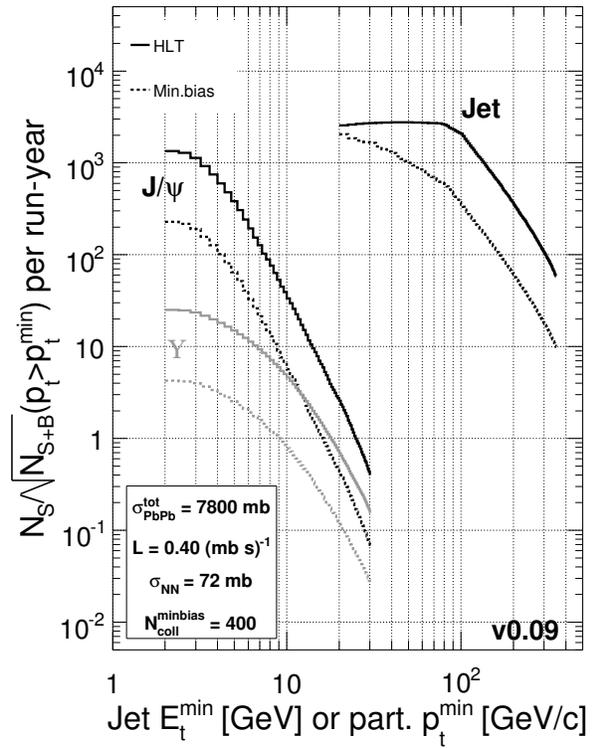


Figure 27: Statistical significance for data taken with the HLT using the trigger settings in Table 7 compared to the minimum bias case.

8.6 Benefits of HLT running at design luminosity

The importance of the HLT is clearly visible in Fig. 26, which compares rates of events to tape with HLT event selection with those for minimum bias running.

Using the HLT, a gain of more than an order of magnitude is achieved for jets at large E_T and for dimuons. Note that for this comparison, the rate for each process was only counted in the corresponding trigger stream, leading to rates below minimum bias for low E_T in the jet channel. At these low E_T , measurements will likely be limited by systematic errors. The advantage of using the HLT is further illustrated in Fig. 27, which shows the statistical significance, $S/\sqrt{S+B}$, as a function of E_T (p_T). A gain in data significance of more than a factor of 5 can be achieved in these channels. Correspondingly, the usable range in E_T (p_T) for the jet and dimuon measurements is extended by more than a factor of 2 and 3, respectively.

In terms of the accessible physics, the benefits of the HLT can be summarized as follows

- Detailed studies of the Υ family will only be possible using the HLT selection. In particular, measurements of Υ' relative Υ yields, which are crucial for studies of the initial temperature of the medium, are only made possible by the HLT.
- Using the HLT, J/ψ statistics are improved such that detailed measurements of the yields as a function of centrality or reaction plane become possible. This will provide information on their interaction with the medium that a measurement without the HLT could not deliver. Another possible example is a study of rapidity distributions of the J/ψ , which could point to the importance of recombination processes in the observed J/ψ yields.
- Statistics and reach for high E_T jet measurements above 100 GeV are greatly enhanced by the HLT. These well defined jets provide a qualitatively new tool in understanding the transport properties of the medium, giving the largest lever arm for testing different models of the interaction of the fast parton with the medium. Experience in p+p collisions and our own model studies show that a large statistics set of jet events and highly differential studies will be necessary to calibrate these new experimental tools. The combination of tracking information with the electromagnetic and hadronic calorimeters, over the large CMS acceptance, in combination with the large statistics allowed by the HLT will be crucial for developing a complete understanding of these new experimental tools.
- Although studies are still ongoing, the production rates for Z^0 production or composite channels like Z^0 -jet correlations show that these measurements will only be feasible using an efficient trigger. Composite channels that allow a clear tag of the jet energy will be critical in calibrating the jet finding performance in p+p or peripheral Pb+Pb events and for obtaining a qualitatively new handle on partonic energy loss in more central Pb+Pb collisions.

8.7 Low Luminosity Scenario

The trigger simulations were repeated for an initial LHC Pb+Pb run projected to be run at 1/20th of design luminosity. We assumed 10^6 sec total run time and an event size 1MByte larger than for later running periods. For the initial run, only 50% of the DAQ event builder CPUs and 50% of the HLT event filter will be installed (see Section 6). Compared to running at design luminosity, the lower initial luminosity will ease the bandwidth and computing requirements, leaving the output bandwidth of 225 MByte/sec as the main constraint. Figures 28 and 29 show the rate and statistical significance results from the low luminosity simulations for minimum bias running and for HLT running using the trigger allocation given in Table 8. Although the lower collision rate limits the gain from using the HLT, the integrated yields are higher than for minimum bias triggering by a factor of 2 to 3 for jet and dimuon triggers, with the biggest gains seen for Υ measurements.

Channel	Threshold	Scale-down	Bandwidth [MByte/sec]	Event size [MB]
Min. bias	–	1	146.25 (65%)	3.5
jet	50 GeV	25	45 (20%)	6.4
J/ψ	0 GeV/c	1	11.25 (5%)	5.9
Υ	0 GeV/c	1	2.25 (1%)	5.9
γ^{prompt}	10 GeV	1	18 (8%)	6.8
UPC/forward coll.	–	1	2.25 (1%)	2

Table 8: Strawman trigger table for running at 1/20th of design luminosity, assigning fractions of the total bandwidth (225 MByte/sec) to individual trigger channels. The last column shows the average event size for each of the trigger streams.

8.8 Summary

We have performed simulations of the HLT performance for studies of dimuons and jets in two luminosity scenarios. These studies validate our trigger strategy to only perform Pb+Pb event rejection in the HLT, based on

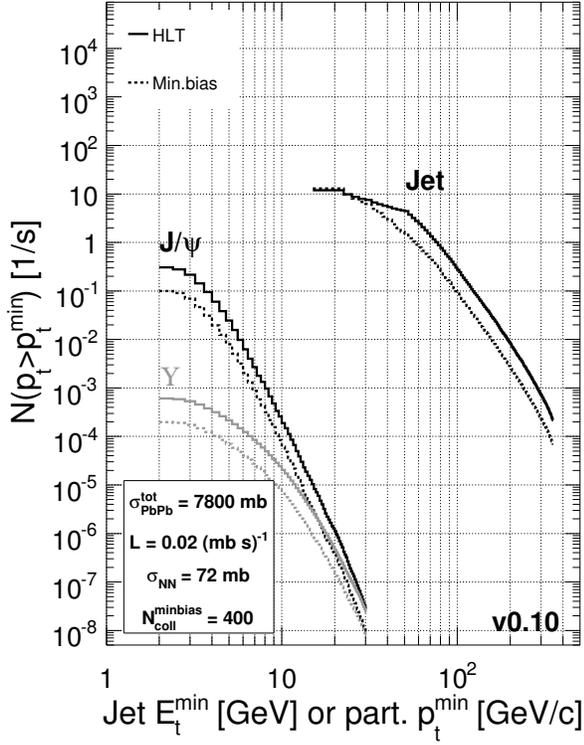


Figure 28: HLT J/ψ , Υ , and jet trigger rates for the trigger settings in Table 8 compared to minimum bias rates for the low luminosity scenario.

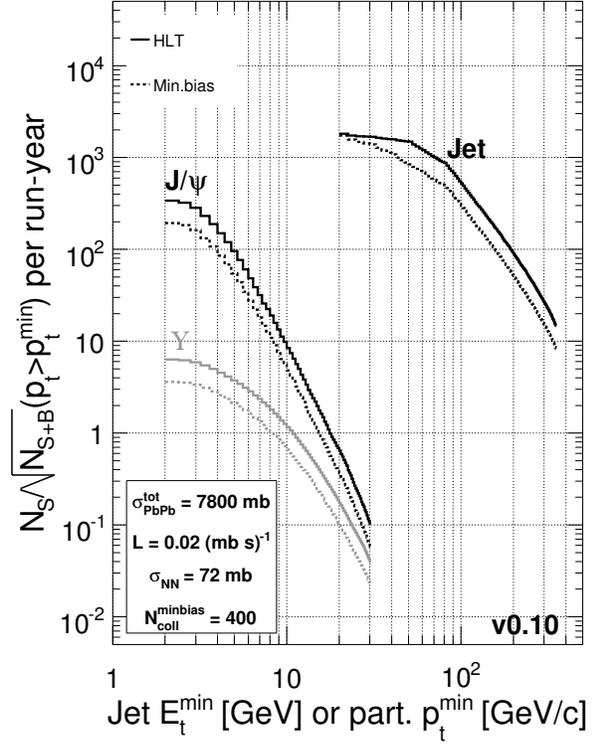


Figure 29: Statistical significance for the case of data taken with the HLT for the trigger settings in Table 8 compared to minimum bias for the low luminosity scenario.

the outcome of running reconstruction algorithms on the full event information. The simulation is based on parameterizations of event-size, timing performance, and algorithmic efficiencies from full event simulations. The results show gains of about a factor of 3 for low-luminosity running and more than an order of magnitude for high luminosity running for the selected trigger channels. The timing studies indicate that the jet finding and $L2$ muon algorithms in their present form use up to 25% of the HLT CPU budget, depending on luminosity. Further optimization of the $L3$ muon algorithm is required to fit within the CPU budget. Additional algorithms for other trigger channels involving electron or photon identification are under investigation and will require a significant fraction of the CPU budget. The studies therefore indicate that the full capacity of the planned HLT filter farm will be needed to maximize the trigger selectivity while maintaining sufficient throughput. A comparison of the physics output rates of minimum bias and HLT running at design luminosity illustrates the crucial role the HLT plays for the CMS heavy ion physics program. The added statistical power provided by the HLT is essential for differential studies connecting rare probes to the physics of the QCD medium.

9 Funding Request

Extensive development of the High Level Trigger is necessary in order for CMS to fully exploit its potential for heavy ion physics. Therefore, we propose that the US CMS heavy ion group focus its effort and resources on the development of the HLT and the software and hardware infrastructure of the Filter Farms. In addition, we propose to contribute to the experiment by purchasing, testing, installing, and commissioning Filter Units in time for the first heavy ion collisions at the LHC.

The implementation of the HLT for the CMS experiment requires participation of heavy ion physicists in the following areas:

1. Verification and development of zero suppression or data compression techniques that might be necessary for the heavy ion data.
2. Implementation of online calibration procedures in coordination with the p+p program, to assure the quality of the online reconstruction.
3. Development of online Quality Control tools to monitor the quality of the heavy ion data.
4. Development of fast physics algorithms to handle the stream of data and reduce it to volumes acceptable for permanent storage.
5. Preparation and familiarization with the software infrastructure of the Filter Farms.
6. Selection, purchase, testing, installation, and commissioning of the Filter Units.

The work requires physics and software expertise, as well as experience in managing online computer farms. The development of online algorithms will require detailed simulations of detector performance with varying assumptions about event rates and backgrounds. The algorithms will need to be tested early on, as a part of the complete DAQ system of CMS. In contrast, the hardware purchases will be organized to profit from the continuing rapid advances in computer technology and will, therefore, occur as late as is reasonable. The proposed budget attempts to balance these two requirements by providing for early resources needed to develop software followed by larger purchases of equipment only a few months before the arrival of the first heavy ion beams at the LHC. Further purchases are planned to be commensurate with the expected growth of luminosity as well as the need to stage the growth of the computer farm in order to avoid large jumps in system complexity.

The Filter Farm is assembled using commercially available components. The estimate of funds needed to provide the necessary elements is based on the recent history of the cost of computer products. Typically, the price of a functional unit, for example a single PC, is approximately constant but the processing power, storage, and network speed grow significantly with the progress of technology. At the present time (May 2006), the estimated price of one PC with 1.8 GHz Opteron dual-core CPU including two Gigabit local network interfaces and 300 GB disk is about 3.1 k\$. We used the empirical “Moore’s Law” to scale expected processor speed and disk storage volume and assume that by the time the LHC starts the network bandwidth will be sufficient to handle the expected data volume. The estimates used to calculate the cost of individual components were done using those presently available. As discussed in Section 8.2 the assumed performance extrapolations result in capabilities which match the requirements by the time of their purchase. Initial trigger studies have identified a number of channels of particular interest to the physics program (see Table 7) and have studied in detail the timing requirements of two of them. These were found to consume about 10–20% of the proposed CPU capacity. A very preliminary version of one of the more sophisticated proposed triggers actually exceeds the planned capacity, indicating the necessity of both additional software development work as well as the full complement of CPUs in the proposed HLT. To achieve the required performance, the proposed purchasing plan depends on a strategy of buying hardware as late as possible, especially for CPU units. If the achieved performance exceeds the actual demands, we will be able to add functionality to the online programming, thereby increasing the complexity of the online algorithms.

The development of sophisticated real-time algorithms requires long, detailed, and careful testing and debugging. The algorithms need to be run in a realistic environment. Hardware purchased relatively early in the project will be used to establish a test bed for various aspects of the software and data handling. We will need the uninterrupted availability of a reasonably-sized test bed to allow us to continue the development and testing of algorithms and

to respond to the possibility of a changing environment at the accelerator. As one illustration of the need for such a facility, the PHOBOS experiment at RHIC suffered from the lack of a separate DAQ development platform. The difficulty of testing DAQ software during data taking resulted in significant delays in improvements to the capabilities of the system. We expect to maintain a development system at MIT throughout at least the first few years of running.

The heavy ion groups will concentrate on expanding the total processing power of the HLT while the CERN and US HEP groups will build the infrastructure of the complete readout and event building system in addition to providing the remaining CPU. This division of tasks will allow very efficient use of the heavy ion efforts in DAQ commissioning, operations, and upgrades. The schedule of purchases and installation will be closely matched to the activities in the experiment and coordinated by the CERN DAQ group under the leadership of Sergio Cittolin. The timeline for the installation is a best estimate at this time. Because of the modular nature of the system, it is not necessary that the hardware be identical in each of the subfarms. In fact, with the continuing advances in computer technology, it is almost guaranteed that there will be differences. Using the test facility at MIT, and in coordination with the DAQ group at CERN, the heavy ion groups will ensure that the hardware to be purchased is fully compatible with the requirements of the HLT.

We plan to purchase the computers in the US and ship them directly to CERN, where they will be extensively tested using standard CMS “burn-in” procedures. We will use all available funding to buy DAQ PCs with the exception of the initial purchase of a small suite of machines which will be used for development and testing of HLT code. The price estimates used in this budget are consistent with the estimates done by the CMS DAQ group for other parts of the Event Builder. The uncertainties in the estimates of required CPU capacity and future CPU capabilities are large, so it makes no sense to introduce financial contingency to the budget calculations. Instead, we assume that as more CPU is available, we will be able to do better physics, both faster and more efficiently.

During the first four years of the US CMS heavy ion DAQ-related efforts, we plan to do the following:

- FY2007: After testing and evaluation of machines from several vendors, we will purchase 30 PCs in April, 2007, ship them to MIT, and build a system with the full functionality of the smallest element of the Filter Farm. Siting of this facility at MIT will allow uninterrupted R&D of HLT algorithms while the DAQ system at CERN is being commissioned for the p+p program. The machines will be mounted in a rack which will be an exact copy of the equivalent racks at CERN such that the ongoing improvements can be easily transferred between the two systems. All members of the US groups will have access to this facility.
- FY2008: Select the best available hardware by using the MIT development system to do additional extensive testing and evaluation of newly available PCs from several vendors. Purchase 200 PCs in April, 2008, and ship them directly to CERN. Conduct testing and “burn-in” at CERN. Install the PCs into CERN-provided racks and commission them with operating software and HLT algorithms. Continue code development in close coordination with our physics studies. Use the processors in the first lower luminosity Pb+Pb run at the end of calendar year 2008.
- FY2009: Repeat the selection and testing procedure at MIT in early 2009. Purchase the second batch of 200 PCs in April, 2009. Install and fully test them in the DAQ system at CERN during the summer. Working at both CERN and MIT, optimize the software infrastructure and algorithms based on the experience gained from the first heavy ion collision data.
- FY2010: Based on an evaluation of the machines purchased earlier in calendar year 2009, purchase the final batch of 200 PCs in October, 2009, right at the start of the new fiscal year. The goal is to install and test this final complement of computers before the arrival of full luminosity Pb+Pb beams in late calendar 2009 or early 2010. Continue software infrastructure and algorithm development at both CERN and MIT.

This proposed strategy of CPU deployment is well matched to the expected ramp-up of heavy ion beam luminosity at the LHC. The initial significant block of CPUs will be in place for the first Pb+Pb run which is projected to be at about 1/20 of the full design luminosity. In this run, the plan is to concentrate on minimum bias and other high rate event types. The processing power of the partially completed HLT will be needed to ensure the cleanest possible event sample and also for data compression to maximize the event rate to tape. The second major purchase is scheduled to allow some contingency in the plan for installation and testing in time for the first full luminosity

heavy ion beams. In order to maximize the CPU power of the full system, the final large purchase will take place as late as possible before this second run.

The construction of the development system at MIT will require technical help as well as hardware including racks, power strips, and network switches. We estimate that we will need 0.25 FTE technician at MIT in FY2007 to set up the test station and conduct hardware checks. A breakdown of the DAQ-related budget items in each of the fiscal years is shown in Table 9. The additional hardware needed to set up the test facility is included in the “Equip.” column. As equipment, the hardware and PCs will not incur overhead charges. The technician cost includes benefits. The next-to-last column lists administrative allocation and service charges levied by the Laboratory for Nuclear Science.

Year	PC Units	PC Cost	Equip.	Tech. (SWEB)	Overhead	LNS	Total
							Costs in k\$
FY2007	30	93	6.5	21.6	15.3	3.6	140
FY2008	200	620	0	0	0	0	620
FY2009	200	620	0	0	0	0	620
FY2010	200	620	0	0	0	0	620
Total	630	1953	6.5	21.6	15.3	3.6	2000

Table 9: DAQ-related budget request showing number of computers and costs in k\$.

The overall cost of the equipment to be installed as an integral part of the CMS experiment is estimated to be about 1.86 M\$ with an additional 140 k\$ for purchasing and installing the MIT test system. Therefore, the total proposed budget request from US Nuclear Physics funds for CMS DAQ-related activities is estimated to be 2 M\$. This contribution from the US heavy ion groups will constitute about 40% of the CMS computing power for the High Level Trigger and Event Filter of CMS, about 10% of the overall cost of the full DAQ system, and about 0.4% of the total cost of the CMS experiment. It should be noted that the CMS budget for future years includes funding for maintenance and gradual replacement of the full HLT, including the portion initially purchased using US Nuclear Physics funding.

10 Personnel, Schedule and Management Plan

Heavy ion physics is an integral part of the CMS experiment, which will be operated during both the p+p and Pb+Pb beams at the LHC. Physicists in the heavy ion groups are full members of the CMS collaboration and actively participate in the detector construction as well as data analysis for both p+p and Pb+Pb collisions. One specific responsibility of the heavy ion members is preparing the experiment for Pb+Pb data, in particular the trigger and online software. Other more general responsibilities include experiment maintenance, shift taking, data processing, and data analysis. Many high energy physics groups in the US have been participating in CMS for some time under an agreement between the collaboration and the US funding agencies, DoE and NSF. Their participation includes leadership roles in the construction of several sub-detectors as well as in the computing and data analysis projects. The US heavy ion groups will follow a similar path and will be part of both the full CMS and the US CMS collaborations. Michel Della Negra of CERN is the spokesperson of CMS and Harvey Newman of Caltech is the Chair of the US CMS Advisory Board.

Preparations for CMS physics analysis are carried out by separate topical groups within the Physics Reconstruction and Simulation group (PRS) which coordinates the major components of the program. Simulation and reconstruction efforts are organized to address specific physics topics which naturally cross the boundaries of individual CMS detector subsystems. The physics categories within the PRS include Higgs Physics, Standard Model Physics, Supersymmetry, and Heavy Ions. The organizational structure reflects the importance attached to heavy ion physics within CMS. Bolek Wyslouch of the US CMS heavy ion group has been appointed coordinator of the Heavy Ion subgroup which consists of physicists from around the globe with the US group playing a leading role. The prominence of US physicists in this area is illustrated in Fig. 30.

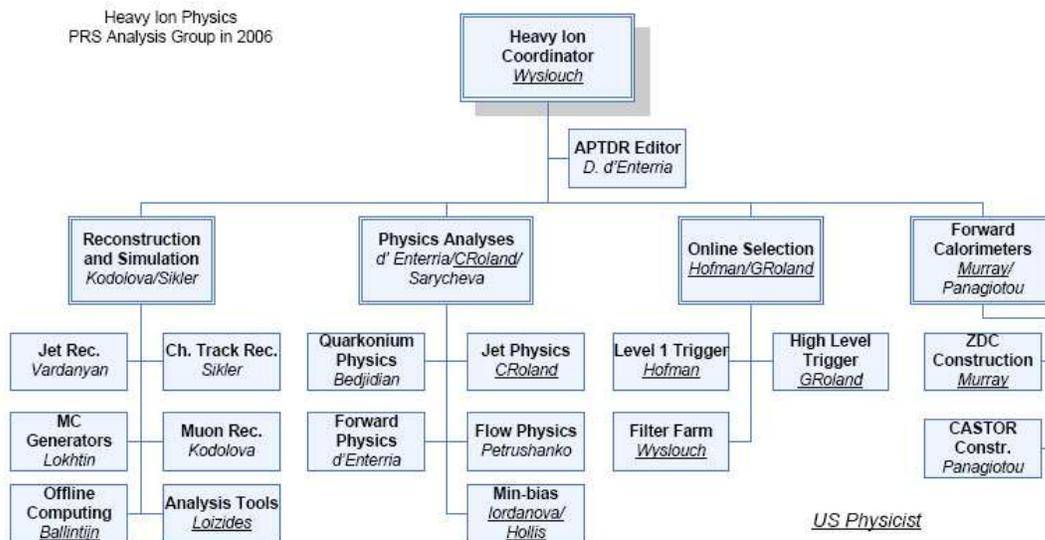


Figure 30: Organization of the heavy ion physics group in CMS. To illustrate their importance in the management structure, the names of US participants are underlined.

In addition, heavy ion physicists are involved in several areas of hardware construction and preparation. In particular, the US groups are building the ZDC calorimeter and part of the DAQ Filter Farm, the latter being the subject of this proposal.

Since the heavy ion physics program is an extension in scope of the US participation in CMS, it needs separate US funding. The overall cost of the CMS Experiment is expected to be about 530 M\$ with the US high energy community contributing about 167 M\$ towards the construction. The expected data taking time during heavy ion runs will be about 10% of the total beam time.

The hardware and operating costs of US heavy ion participation can be separated into five categories (see Section 9 and Appendices B, D, and C for detailed discussion) :

- Experimental equipment to be installed at CERN (ZDC and DAQ Filter Farm, see Section 9 and Appendix D);
- Yearly contributions to the running costs of the experiment which are proportional to the number of participating PhD level physicists (Category A, see Appendix B);
- Maintenance and repair of the installed equipment (Category B, see Appendix B);
- Offline computing resources for data processing and data analysis (see Appendix C);
- Increase in personnel and travel cost for on-site participation at CERN above what was typical for the RHIC experiments (see Appendix B).

The hardware and computing activities will be conducted within the structure of the existing CMS management which is shown in Fig. 31.

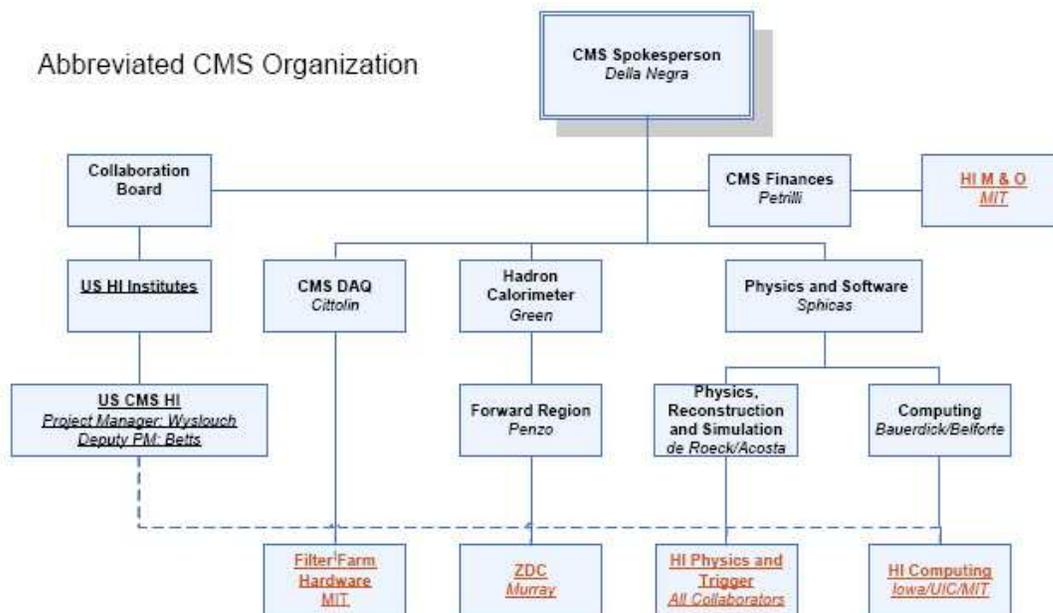


Figure 31: Main areas of activity of the US heavy ion groups within the structure of the CMS management. Note that only the components with US heavy ion involvement are shown.

The Filter Farm purchases and commissioning will be coordinated with the DAQ group led by Sergio Cittolin of CERN. The interests of the US heavy ion groups will be represented by MIT with groups from UC Davis, UIC, Iowa, and Rice participating actively in this effort. The specific division of work will depend on the status of the project and may evolve with time but at this stage the groups plan to concentrate on the following areas of hardware and computing work:

- Development of HLT algorithms for specific physics topics: All groups, depending on physics interest
- Level-1 optimization: UIC
- Benchmark and validation of algorithms in the HLT environment: MIT
- Online data quality monitoring: UC Davis
- Local and remote event data storage management: Iowa
- Communication and control in the HLT farm: Rice

11 Appendices

A Profiles of the collaborating groups

The project described in this proposal will be carried out by the US members of the heavy ion community within the CMS collaboration. The US heavy ion subgroup presently consists of ten institutions. They represent a cross section of the RHIC community and include personnel with wide ranging expertise in heavy ion physics and detector construction, most significantly data acquisition, higher level triggers, and zero degree calorimeters. We estimate that by 2007-2008 the number of US heavy ion participants will total roughly 20 faculty and senior research personnel, 10 postdocs, 20 graduate students, and 10 undergraduate students. The non-US heavy ion contingent in CMS will number about 50 people. The efforts of these nuclear physics (NP) participants is assisted by the large number of high energy physics (HEP) members in CMS, in particular the US physicists at the institutions included in the heavy ion effort. Since many aspects of the data-taking and analysis will be done jointly by the whole collaboration, the productivity of the core heavy ion group will be greatly enhanced by their interaction with these other CMS collaborators. The group members are listed in Table 10. The relevant experience of each group, along with information about related activities and resources at each institution are described below.

A.1 University of California at Davis

The NP group at the University of California at Davis has focused on relativistic heavy ion physics since the early 1990's. They have experience with large acceptance experiments at LBNL's Bevalac (EOS), CERN's SPS (NA49), and BNL's AGS (E895). The experimental group consists of one tenured faculty (Daniel Cebra), one tenure-track faculty (Manuel Calderón de la Barca), and two graduate students with a postdoc search currently ongoing. The analysis of data from STAR is currently the major activity of the group. Daniel Cebra has been involved in the STAR collaboration at RHIC since the original letter of intent and has contributed to many hardware efforts. These included the design and implementation of the FTPCs; the installation, testing and quality assurance to the electronics for the central TPC; and the laser calibration systems for the central and forward TPCs. Manuel Calderón de la Barca has been a member of STAR since 1997 and has participated in several STAR offline software projects, and served as Deputy Offline Reconstruction Leader for STAR for one year. He has been coordinator of the quarkonium triggers program, working on the design and implementation of Level-0 and Level-2 triggers. Ramona Vogt has been a proponent of heavy ion physics at the LHC since the inception of the program. She has contributed to the interpretation of quarkonium and heavy flavor physics in heavy ion collisions at RHIC and the SPS, and will continue to provide theoretical advice on the physics content of the data. The UCD group also has experience building and operating a medium scale computer cluster, experience which will be valuable for the development of online and off-line computing for CMS. In addition, the CMS experiment offers an opportunity for collaboration with the strong HEP group at UCD, which has been active in the US CMS collaboration.

A.2 University of Illinois at Chicago

The University of Illinois at Chicago group has tenured and tenure track faculty in both HEP (Adams, Gerber, Varelas) and NP (Barannikova, Betts, Hofman). The UIC NP group has been involved in heavy ion research for many years, starting with experiment E917 at the AGS and continuing with the RHIC program. The group is currently involved in data analysis for the RHIC STAR experiment as well as finalizing results from the completed PHOBOS experiment. Olga Barannikova is convener of the STAR spectra working group. Russell Betts has previously been co-spokesman and project manager of the APEX experiment and also acted as project manager of PHOBOS during the preparation and startup of the 2001-2002 run. Dave Hofman has previously acted as project manager of PHOBOS during the d+Au run in 2002-2003 and is co-convener of the PHOBOS multiplicity working group. The UIC group had overall responsibility for the design and construction of the PHOBOS multiplicity and vertex detector, and also contributed to the multiplicity analysis algorithms which provided the first physics results from RHIC. Edmundo Garcia was a member of the RHIC group which developed the ZDCs which are common to all four experiments, and is currently contributing actively to the ZDC effort at CMS. Aneta Iordanova and Richard Hollis are post-docs who have been working on the L1 trigger in CMS. In the near future, the UIC group expects to hire another post-doc and continue its participation in the STAR experiment while transitioning all PHOBOS related efforts to the CMS heavy ion program. The group has accepted responsibility for the Level-1 trigger in CMS

for the heavy ion running, and is fully engaged with the LPC at FermiLab and the Wisconsin group, who have the overall p+p trigger responsibilities for CMS. In this regard, UIC has also taken on responsibility for both the p+p and Pb+Pb minimum-bias trigger in L1 at CMS. The group effort will continue to be enhanced by coordination of effort with Fermilab and with the UIC HEP group in CMS, which has been a member of the CMS collaboration since 1995 and has actively contributed to the construction and monitoring of the CMS silicon tracker.

A.3 University of Iowa

The University of Iowa has both HEP and NP groups. The HEP group has been active in the CMS project since its inception. The group originated, designed, and prototyped the CMS-HF detector and is responsible for numerous components of the final detector. The local group also has experience with the simulation software used to study heavy ion collisions in CMS. Yasar Onel is the US Coordinator for HF and the Project Manager for the PMT photodetector project. Edwin Norbeck, who joined CMS in May 2001, has extensive experience in heavy ion nuclear physics and the development of innovative charged-particle detectors.

A.4 University of Kansas

The NP group at the University of Kansas has been a major contributor to the BRAHMS experiment at RHIC, with hardware and analysis responsibility for the BRAHMS multiplicity array, a mid-rapidity threshold Čerenkov detector, and the ZDC detectors. The group currently consists of two faculty (Michael Murray and Steve Sanders), three postdoctoral assistants, and two graduate students. Michael Murray participated in the design, prototyping and integration of the ZDCs used by all the RHIC experiments, built and installed the ZDCs for the BRAHMS experiment, and currently has management responsibility for the CMS ZDC detector development. Steve Sanders was responsible for the design and construction of the BRAHMS Multiplicity Array and mid-rapidity threshold Čerenkov detector. He is currently working on BRAHMS analysis with responsibility for the azimuthal flow, global particle production, and centrality measurements. BRAHMS is expected to take its last data this spring and the analysis is expected to be completed in 12-18 months. During the next year Steve Sanders will be shifting his research focus over to CMS. Oleg Grachov has worked on electromagnetic calorimeters in experiments at Fermilab as well as the STAR experiment at RHIC. The personnel that will participate in the CMS heavy ion program include both faculty members, Oleg Grachov, one or two graduate students, and several undergraduate students. Oleg Grachov is currently working full-time on the ZDC development for CMS. The KU HEP group is working on the CMS Si tracking detectors, and a mutual exchange of information and advice between the NP and HEP groups can be expected.

A.5 Los Alamos National Laboratory

The focus of the Los Alamos High Energy Nuclear Physics team in recent years has been the construction of the muon tracker for PHENIX and the analysis of J/ψ related physics of the QGP. Gerd Kunde recently moved to Los Alamos from Yale, where he was STAR co-convenor of the high transverse momentum working group and project leader for the Ring Imaging Čerenkov Detector. As a LANL technical staff member, he also holds an adjunct professor position at UIC. As part of its ongoing effort, the group obtained a 3 year Exploratory Research grant under the Los Alamos LDRD program to conduct studies using high p_T jets in the CMS experiment to probe the high density medium. The group currently consists of post-docs Paul Constantin (specialist for jet correlations in PHENIX) and Camelia Mironov from STAR, as well as UIC graduate student Miroslav Mihaylov and post-BS student Maria Castro from Wayne State.

A.6 University of Maryland

The University of Maryland's participation in the heavy ion physics program at CMS is within the research group of Alice Mignerey. This group was a charter member of the PHOBOS collaboration at RHIC and took major responsibility for the design, construction, and commissioning of the trigger system for that experiment. The group has a tradition of research into the dynamics of heavy ion reactions, starting with investigations of damped reactions in the 1980's, to intermediate energy studies of multifragmentation at the Bevalac, to studies of collective flow at the AGS, RHIC, and now the LHC. While the group actually resides in the Department of Chemistry and

Biochemistry, there is a strong component in CMS within the Particle Physics Group in the Maryland Physics Department. Thus, participation in CMS provides the Nuclear Chemistry group with a natural transition to the study of the properties of nuclear matter at very high energy densities.

A.7 Massachusetts Institute of Technology

Members of the Relativistic Heavy Ion Group at MIT have a long standing interest in understanding the mechanism of particle production in high energy collisions, in the phases of QCD matter, and in the properties of matter at extremes of energy and baryon density. Their studies of this physics include experiments E178, E451 and E565 at FNAL; NA35, NA49, and WA98 at CERN; E802, E859, E866, and E917 at the BNL AGS; and PHOBOS at RHIC. The planned initial focus of the MIT group's measurements at the LHC is to a large extent the consequence of the experience and results obtained by the group at lower energy.

The MIT group is presently engaged in analyzing data from the PHOBOS experiment. MIT proposed the experiment and has led the construction and installation efforts. Wit Busza is the spokesman of PHOBOS and Bolek Wyslouch was the project manager responsible for the construction. George Stephans was the installation manager. Gunther Roland is presently the physics coordinator of the experiment. In addition, the group consists of several research scientists with expertise in the fields of silicon detectors, data acquisition, electronics, and computing. The MIT group has significant expertise in jet finding algorithms and experience in studies of the systematics of the production of high p_T particles. Gunther Roland, Maarten Ballintijn, Constantin Loizides, and graduate student Corey Reed have been active in the development of the ROOT and PROOF software packages which are now widely used in analysis in both high energy and heavy ion physics. The transition of the group effort towards the heavy ion program at CMS is well underway. The faculty, research staff, and students have been working on preparation of the algorithms for data analysis and high level trigger to begin the commissioning of the proposed Event Filter Farms. The HEP group at MIT is a major contributor to CMS in the areas of data acquisition and computing. We expect to coordinate the Heavy Ion Event Filter activities with the ongoing HEP DAQ effort. The synergy of HEP and NP efforts will result in a very productive group.

A.8 University of Minnesota

The University of Minnesota is represented by a theorist, Joseph Kapusta. He has been a proponent and contributor to the fields of high energy heavy ion collisions and QCD at high temperature since the 1970's. His role is to advise and work on the physics content of the data collected by the collaboration. He may also have a postdoc and/or graduate student involved in this project. There is a strong HEP component of US CMS at Minnesota, represented notably by Professors Roger Rusack, Yuichi Kubota, Prisca Cushman, and Jeremy Mans, which is very interested in the heavy ion program at the LHC. Synergy between the two groups will result in a very productive effort.

A.9 Rice University

The Bonner Lab at Rice University is home to both HEP and NP groups which have often worked together on projects of mutual interest. The HEP group has been very active in the CMS project, including designing and building the electronics and algorithms for the Level 1 Endcap Muon Trigger. The NP group has been involved in STAR from the beginning of the experiment. The Central Trigger Barrel was built at Rice and the group has been leading the effort to develop and construct a time-of-flight detector for the experiment. Moreover, Rice played a key role in the development and implementation of the very successful high level trigger in STAR. Taking into account the involvement of the HEP group in CMS, it was only natural that the interests of the NP group would transition to heavy ion collisions at higher energies.

A.10 Vanderbilt University

Vanderbilt University has a strong program in both NP and HEP. The HEP group joined CMS in 2005 and is involved in both trigger/DAQ development and offline computing. The NP group has a long and continuing involvement with the PHENIX experiment at RHIC as well as prior relativistic heavy ion experience with BNL AGS experiments E814 and E864. The group currently has three faculty members (Greene, Maguire, and Velkovska),

two post-docs, and four graduate students and is in the process of hiring a third post-doc. The group has made vital contributions to PHENIX in every aspect of the experiment: hardware, software, physics analyses, and paper preparation. Prof. Greene is responsible for the PHENIX pad chambers, Prof. Maguire is the simulations manager, and Prof. Velkovska has built an upgrade time-of-flight detector for PHENIX and has been a convener of the PHENIX hadron physics working group. Vanderbilt has significant presence in the PHENIX detector and executive councils and played a major role in preparation of the PHENIX White Paper. The group also has experience with large scale off-line computing. The topics of the group's major physics interests have natural continuations to the LHC energies. The transition from PHENIX to CMS will happen over several years starting with 1/2 post-doc involvement during 2006 and gradually increasing to two students and two post-docs by 2009. During this period, the three faculty members will ramp up to a significant involvement in CMS while continuing to fulfill their responsibilities to PHENIX.

A.11 Acknowledgment of non-US Participants

We are indebted to many members of CMS collaboration working on preparation of both the p+p and heavy ion physics program. In particular Marc Bedjidjan, Olga Kodolova, and David d'Enterria were very helpful in preparation of this proposal. We also need to thank the CMS and US CMS management for their constant support and encouragement, among them we are thankful in particular to Michel Della Negra, Dan Green, Bob Cousins, Harvey Newman, Tejinder Virdee, Daniel Denegri, Sergio Cittolin, and Paris Sphicas.

Institution	Name	Position
U. California at Davis	M. Calderón de la Barca Sánchez	Faculty
	D. Cebra	Faculty
	R. Vogt	Adjunct Faculty*
U. Illinois at Chicago	O. Barannikova	Faculty
	R.R. Betts	Faculty
	E. García	Research Faculty
	D. Hofman	Faculty
	R. Hollis	Post-doctoral
	A. Iordanova	Post-doctoral
U. Iowa	N. George	Adjunct Faculty
	F.D. Ingram	Adjunct Faculty
	E. Norbeck	Faculty
	Y. Onel	Faculty
U. Kansas	O. Grachov	Post-doctoral
	M. Murray	Faculty
	S.J. Sanders	Faculty
Los Alamos Nat. Lab.	P. Constantin	Post-doctoral
	G.J. Kunde	Research Staff
	C. Mironov	Post-doctoral
U. Maryland	A.C. Mignerey	Faculty
	M.B. Tonjes	Post-doctoral
Mass. Inst. of Tech.	M. Ballintijn	Research Scientist
	W. Busza	Faculty
	C. Loizides	Post-doctoral
	C. Roland	Research Scientist
	G. Roland	Faculty
	G.S.F Stephans	Senior Research Scientist
	B. Wyslouch	Faculty
U. Minnesota	J. Kapusta	Faculty*
Rice U.	B. E. Bonner	Faculty
	G. Eppley	Research Staff
	W. Llope	Research Faculty
	G. Mutchler	Faculty
	P. Yepes	Research Faculty
Vanderbilt U.	S.V. Greene	Faculty
	C.F. Maguire	Faculty
	J. Velkovska	Faculty

Table 10: Current institutions and non-student members of the US component of the CMS heavy ion group. Names of theorists participating in the program are marked with *.

B Operations Budget

B.1 Collaboration Costs

The CMS collaboration, following the long established CERN practice, maintains a fund to which all collaborating institutions or countries contribute. This fund is used to provide resources for detector operations at CERN. Examples include gas supplies for drift chambers, maintenance contracts for installed equipment, maintenance technicians, and replacement costs of outdated electronics. In particular, the computers used for data acquisition will be gradually replaced and upgraded. This continual replacement of the online CPUs will significantly enhance the system's capability by taking advantage of the expected increases in CPU speed and disk space density. The size of the contribution to this "M&O Category-A" fund is calculated annually, each September, and is proportional to the number of active PhD physicists in the group. It is important to note that a person needs to be a collaboration member in full standing for at least one year in order to appear as an author on CMS physics publications. Table 11 gives a present estimate of the number of the US physicists expected to work in the CMS heavy ion program. The contributions for the physicists working on the heavy ion program will need to be covered by heavy ion physics funding agencies. It is expected that the contribution per physicist in FY2007 will be about 7.7 k\$ without overhead, increasing gradually to about 10.5 k\$ for FY2010 and beyond. "M&O Category-B" funds cover the responsibility of groups that have built detectors to maintain their equipment. Currently, the only item in this category for the US CMS heavy ion group is the Zero Degree Calorimeter which will need yearly maintenance with a likely replacement of the tungsten and fibers in FY2010 or FY2011, after a few years of operation. Note that upgrades and maintenance of online computing will be covered using "Category A" funds collected from the entire CMS collaboration.

B.2 Additional Travel Costs for US Institutions

The large distance to CERN will result in a different mode of operation for the participating US groups compared to research at RHIC. Multiple short trips will be replaced by less frequent, and longer, stays at CERN. This will result in some increase in travel costs which will vary from group to group. The role of videoconferencing and remote collaboration tools will become more important. The presence of the CMS Virtual Control Room and Physics Analysis Center at Fermilab will further reduce the need for trips to CERN. To provide the essential direct contact with experts at CERN, we expect that typically 1-2 physicists from each institution as well as the graduate students will stay at CERN for extended periods of time, totaling roughly 14 FTEs on site. Adjustments for cost of living and currency exchange protection are expected to increase costs about 5–15 k\$ annually per person (lower for students than for PhDs) stationed full time at CERN compared to what was typical at Brookhaven. For the entire US CMS heavy ion group of about 30 PhDs with an additional 15-20 graduate students, we expect that the overall cost of travel will be about 250 k\$ more than the cost of the same groups doing research at RHIC.

The total operating budget without overhead for the US CMS heavy ion program, including collaboration fees and the expected increase in travel costs is summarized in Table 11.

Year	PhDs	Students	M&O Cat A	M&O Cat B	Travel	Total
	Total number		Costs in k\$			
FY2007	20	10	154	5	200	359
FY2008	30	18	276	5	250	531
FY2009	30	18	306	5	250	561
FY2010	30	18	321	5	250	576
FY2011	30	18	321	30	250	576

Table 11: Increase in operations costs without overhead for the US CMS heavy ion program in k\$.

The operations of most of the groups involved in this proposal are currently funded by the DoE Office of Nuclear Physics. The group at the University of California at Davis is funded by the National Science Foundation. The Los Alamos National Laboratory group is currently funded by the lab's LDRD program and is seeking DoE Nuclear Physics funding to continue their involvement afterwards. The University of Iowa group is also seeking DoE Nuclear Physics funding to cover its CMS heavy ion research program.

C Computing for Heavy Ions

The general CMS Computing Model is described in detail in “The Computing Project Technical Design Report” [53], which introduces the event model, the distributed and layered computing infrastructure, and the data flow. That document focuses primarily on the requirements of the proton-proton running but also takes into account infrastructure for collecting and archiving the heavy ion data on site at CERN as well as the initial processing in Tier-0. Resources are also available for transporting this data to the Fermilab facility and archiving it there, but the cost of tapes at Fermilab is not covered. At present, off-line computing resources outside CERN have not been specifically allocated to heavy ion analysis. In this appendix the general model is briefly presented, followed by an overview of the specific requirements for computing and data management with heavy ion data.

C.1 The CMS Computing Model

The CMS computing model defines a hierarchy of computing centers located around the world that will provide both data storage and processing for the experiment. The centers in each layer or tier have a well defined set of responsibilities. Technologically advanced computer networks are used to transfer the data between the centers.

Within the model there are four major types of event data. The “RAW” type encompasses objects created by the HLT from the detector data. The “RECO” type contains the data produced by the event reconstruction software and includes objects like tracks, vertices, etc. including full, detailed information on more basic data such as hits. The “AOD” (Analysis Object Data) is derived from the “RECO” information and provides data for physics analysis in a convenient, compact form. The “TAG” type provides a highly condensed summary of each event and is meant for rapid identification of events for further study. In addition there will be a “MC” type for storing simulated events.

At the top of the hierarchy of computer centers is the Tier-0, located at CERN. It receives the raw data from the detector during running and is responsible for archiving the primary copy of the data to tape. The first pass reconstruction of the events is performed in real time at the large computer farm of the Tier-0. In parallel, the raw data is distributed to the next layer, the small number of Tier-1 centers located in different countries. These centers are responsible for storing a second copy of the raw data. Tier-1 centers also supply processing resources for further reconstruction, calibration, and other data intensive analysis tasks. It should be noted that the costs associated with archiving a second copy of the heavy ion data at the Tier-1 level are not included in the current CMS plans.

Batch processing of data analysis as well as Monte Carlo simulations are performed at the larger number of Tier-2 centers which each provide a significant amount of CPU power and disk storage. Data required for these end user analyses is staged from the Tier-1s. The Tier-2s have no custodial responsibility for data, but function as large disk caches. Finally, the Tier-3 centers provide interactive access to local groups and additional modest computing capacity for the experiment. Users at the Tier-3 centers will rely on Tier-2 for access to the data. Again, it is important to note that the CMS planning makes no provision for analysis resources allocated to heavy ion physics.

C.2 Offline Computing for Heavy Ion Runs

The computing model for heavy ion data closely resembles the general model described above. The goal is to use the standard CMS infrastructure for transferring, storing, and accessing the data, but using a subset of the general resources as well as resources dedicated to heavy ion analysis. The standard CMS software framework for reconstruction and analysis will be employed. The heavy ion groups have already started participating in software development, including the critical validation for heavy ion events which are typically much larger in size. It is planned to incorporate dedicated, heavy ion specific algorithms and optimizations where needed.

However, the heavy ion data and analysis are sufficiently different that adjustments must be made. As such, it makes sense to place all of the heavy ion data in one or more Tier-2 or comparable centers. Possible candidates for these sites currently include Iowa and Vanderbilt, which have presented bids for hosting Tier-2 centers; MIT, where the high energy and heavy ion groups are already cooperating in the installation of a Tier-2 facility; and UIC, which is seeking state funding to establish a dedicated CMS heavy ion computer facility. The plans presented here are based on the provisional run plan included in Table 2.1 of [53]. The dominant constraint is the available tape writing speed at the Tier-0, currently set at 225 Mbyte/sec. The goal is to have a system which is sufficient to process and analyze the collection of data that will accumulate over the consecutive data taking periods.

C.2.1 Heavy Ion Data Flow

The offline part of the data processing will start with a small number of streams generated by the HLT farm, using tags based on various triggers. The actual number of streams is likely to be less than is planned for p+p data. It is foreseen that the computing resources at the Tier-0 will not be sufficient to process all data so only certain streams will be selected for prompt reconstruction, based on physics priorities and the available capacity. Both the reconstructed and unreconstructed streams will be recorded to tape at the Tier-0 and also transferred to the Fermilab Tier-1. The streams which were not yet reconstructed will be processed at the dedicated Tier-2s and the reconstructed events will be sent back to Fermilab for tape archiving.

C.2.2 Heavy Ion Data Storage Requirements

Storage on tape at Tier-0 and Tier-1 will provide reliable long term archiving, and semi-permanent disk caches, mostly at the Tier-2 level, will provide access to events for reconstruction and analysis jobs. Several of the planned heavy ion analyses will likely require access to tracker hits at the RECO level. Therefore, it is expected that, in addition to the AOD data, all, or at least a large fraction of, the RECO data will need to be available at the heavy ion Tier-2s. Based on the planned data-taking capacity, data volumes are expected to be ~ 250 TByte of RAW data, which probably increases to ~ 430 TBytes of RECO data for each full luminosity heavy ion run. Based on a quote from the Fermilab Tier-1 facility, the cost of tape storage of the heavy ion data, as well as Monte Carlo simulation results, is expected to be 75 k\$ annually starting in FY2008.

C.2.3 Heavy Ion Computing Requirements

The computing required for processing heavy ion data can be split in three categories; reconstruction, end user analysis, and Monte Carlo simulation. In heavy ion events, the time needed for reconstruction varies greatly with the particle multiplicity. The distribution of multiplicities in an event is a function of the physics as well as the selection made by the HLT. It is also foreseen that the offline reconstruction will further select only a fraction of events for which full track reconstruction, which requires the most CPU time, will be performed. Over time the understanding of the input parameters for the reconstruction software, such as alignment and calibration, will improve. The reconstruction algorithms will also be improved, or replaced by better ones. Therefore, it is expected that the data will have to be reprocessed, probably multiple times in the first year(s).

The physics analysis of heavy ion data will require access to both RECO and AOD data depending on the type of analysis. The one or more heavy ion computing centers will need sufficient CPU capacity for these additional tasks, which, based on previous experience, might be as much as 50% of the amount needed for reconstruction.

Simulated event data, created using Monte Carlo techniques, play a crucial role in basic studies of the detector response, the development of reconstruction software, and the comparison to predictions of physics models. Generating these simulated events is very CPU intensive, but does not typically require a large amount of I/O, so this task is very well suited for use of Grid resources. The amount of simulated events required for the physics analyses and detector studies is not known in advance. Based on the experiences of previous heavy ion and hadron collider experiments, it is foreseen that an event total on the order of 5% of the real data will need to be generated.

C.3 Heavy Ion Computing Purchase Plan

It is expected that significant computing resources beyond those allocated in the current CMS plan will be needed to carry out off-line data reconstruction and analysis for heavy ion physics. The US groups would like to store all raw heavy ion data locally in order to facilitate full reconstruction and analysis by the faculty and students. Based on experience with the RHIC experiments, it is desirable to have sufficient local storage for the processed data from at least two years of running, in particular to allow efficient analysis of rare events. In order to keep the analysis commensurate with the running schedule, it is necessary to have sufficient processing capacity to fully reconstruct a significant fraction of the data from one run within the one year expected between runs. The future capability of available CPUs is not well known but can be estimated by projecting the recent growth per year of roughly 45–50% in CPU speed and disk size for constant cost. Clearly the final algorithms are still being developed but sufficient software infrastructure exists to estimate analysis times on current machines. Using this information, the resources provided by a staged installation can be projected. It is estimated that such a total facility (perhaps

divided into two or more locations at the collaborating universities) would provide the required resources, i. e. enough disk space for simultaneous local access to several years worth of data and enough processing speed to fully reconstruct a significant fraction of all heavy ion events. The cost of providing these critically necessary resources is estimated at 325 k\$ annually. Some of the details are given in Table 12. Note that in FY2007, only simulation results will need to be stored, and therefore tape costs are smaller. As a consequence, more money is dedicated to buying CPUs. In each of the following years, a roughly equal number of CPUs will be purchased, while machines that are older than 3 years will be replaced. Therefore, the maximum physical size of the facility will be reached in FY2010, but the CPU and disk capacities will continue to expand. The total heavy ion computing cost is 400 k\$ per year, starting in FY2007 and continuing for the duration of LHC running. This sum includes a small charge for tape storage at the Tier-1 center at Fermilab in FY2007 increasing to 75 k\$ annually thereafter. As capital equipment, these charges would not incur overhead. Note that 100 k\$ of the 400 k\$ in computing purchases listed for FY2007 will be covered using funds granted in FY2006 in response to a supplementary request from the MIT group. Therefore, the FY2007 computing costs will total 300 k\$. The FY2007 entry in Table 12 includes items to be purchased using both FY2006 and FY2007 money.

This budget includes only the tape storage and computing costs and does not cover infrastructure such as power and cooling. Those would need to be provided by the host institution(s) or other funding sources. Possibilities currently being actively pursued include sharing infrastructure with CMS high energy Tier-2 facilities under construction at MIT and proposed at Iowa and Vanderbilt, as well as using matching funds from the State of Illinois towards the construction of a CMS heavy ion computer facility at UIC. Multiple locations can be considered since the cost for the required computer resources is expected to be the same whether installed at one or more sites. The decision on siting the heavy ion computing is expected to be based largely on infrastructure availability.

There are good reasons to construct the computing centers in several yearly increments. This allows one to benefit from the growth of computing power at a fixed cost per machine. From a practical point of view, gradual staging makes the total deployment easier by having a number of more manageable steps rather than a single big deployment. Risk is reduced by having multiple purchases, possibly from different vendors. Finally, the requirements for heavy ion data storage, reconstruction, and analysis will grow over time, further supporting this model.

Fiscal Year	Tape [TB]	Batch Nodes	CPU [SI2k.s.10 ¹²]	CPU Total [SI2k.s.10 ¹²]	Disk [TB]	Total Disk [TB]
2007*	9	105	20	20	157	157
2008	430	83	23	43	186	343
2009	430	83	33	75	279	623
2010	430	83	47	103	419	884
2011	430	83	69	149	628	1326

Table 12: Computing resource purchase schedule. * Note that the FY2007 entry also includes computers purchased in FY2006. See text for details.

Based on these preliminary estimates, such a setup will support the basic storage and analyses needs expected for the heavy ion program. Depending on the availability of extra resources, and on software improvements, it would be possible to perform even more advanced analyses. The participating groups will continue to work both individually and collectively to explore the availability of additional resources for off-line computing. Additional capacity, especially in the area of CPU power, would allow more flexibility in the scheduling of analyses, more opportunity to explore a broader variety of physics topics, and an improved ability to respond expeditiously to new discoveries.

D Zero Degree Calorimeter

The US CMS heavy ion group will provide a pair of Zero Degree Calorimeters (ZDC) which will serve a vital role in the heavy ion program at CMS. The ZDCs are mounted behind the first accelerator magnet at the end of the straight section surrounding the interaction point. Located 140 m from the vertex and installed into a space 1 m long and 9.6 cm wide between the two beam pipes, they cover $|\eta| \geq 8.5$ for neutral particles. This detector is incorporated into the TAN, the structure which protects the accelerator magnets from synchrotron radiation. The calorimeter design is similar to that for the RHIC ZDCs [50] but with the addition of a segmented electromagnetic section which allows the detection of photons as well as neutrons. The Kansas, Iowa, and UIC groups are primarily responsible for the two ZDCs.

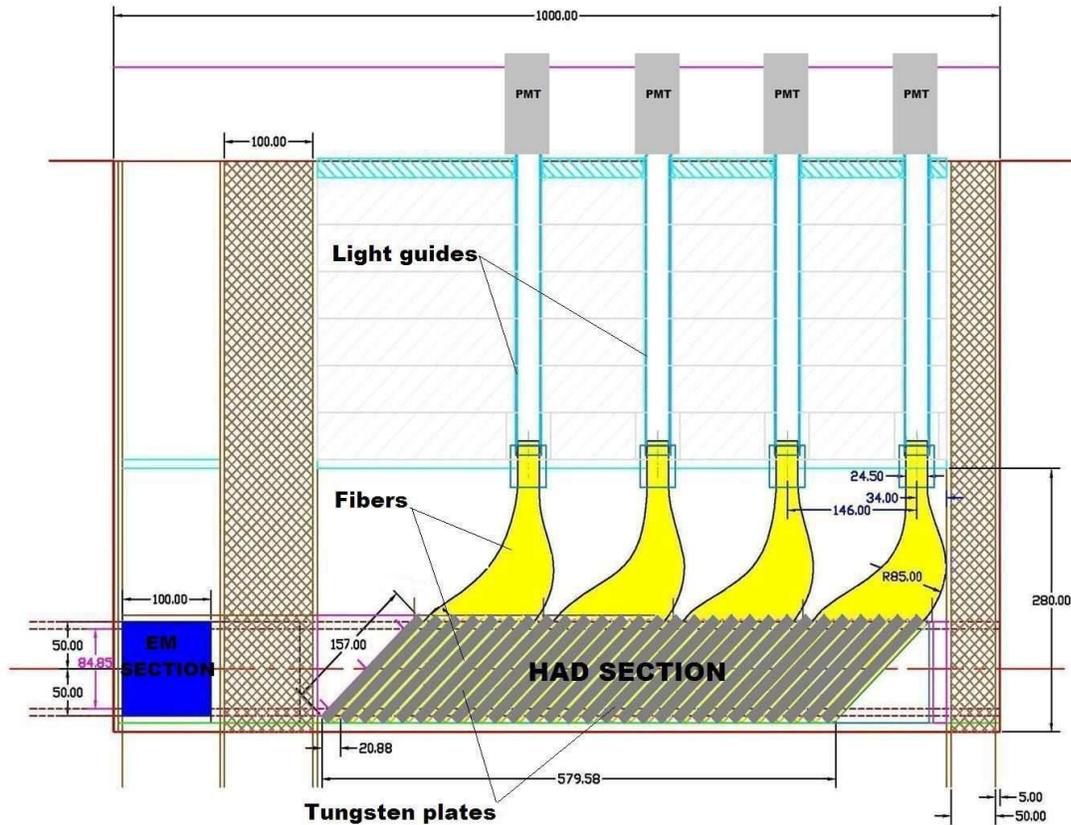


Figure 32: Side view of the ZDC showing the electromagnetic and hadronic sections. A luminosity monitor will be mounted in the space in between. Dimensions shown are in mm.

Figure 32 shows a side view of the calorimeter with the electromagnetic (EM) section in front and the hadronic section behind. A luminosity monitor being built by the LBL Accelerator Physics group will be mounted in between the two. A 5 cm space is reserved for a proposed flow detector. The ZDCs are Čerenkov calorimeters which sample only the high energy core of the induced showers. The hadronic part is segmented into 4 longitudinal sections composed of alternating 1.5 cm thick tungsten plates mounted at a 45° angle as shown with quartz fibers in between. In the EM section, the plates are 2 mm thick and mounted vertically. The use of tungsten provides sufficient density to contain the transverse spread of the signal to within a few cm of the initial trajectory. Quartz fibers provide high radiation tolerance. The fibers in the EM section are divided horizontally into 5 segments, giving the ability to measure the location of single photons to a precision of a few mm. By determining the average location of photons emitted roughly collinear with the beam, the ZDC will be able to monitor the beam-crossing angle. The phototubes (PMT) will be shielded from albedo by a matrix of boronated-plastic (to cut down neutrons) and lead (to absorb photons). The PMTs are 8 stage Hamamatsu R7525HAs which are optimized to match with

the CMS readout system [51] which uses differential input to reduce noise. A large dynamic range is accomplished through a multi-range technique. The Instrument Design Lab at KU and FNAL are designing a system to bleed off $\sim 5\%$ of the signal for timing information which will determine the vertex to a resolution of 3 cm for use in the Level 1 trigger. In addition to the physics capabilities discussed below, the ZDCs will also serve a vital function in the tuning and monitoring of the heavy ion beams, a capability which has been well established at RHIC. Several communication channels linking CMS to the LHC control room have been reserved for real time ZDC data on coincidence and background rates as well as beam-crossing angle.

D.1 Physics Capabilities of the ZDC

The ZDC signals can be used for a variety of physics studies, some of which are well established at RHIC and some of which are unique to the higher LHC energies. The ZDC contributions to the main heavy ion physics program are primarily in the area of event characterization, namely centrality determination and flow.

In any study of heavy ion collisions, the variation in impact parameter either needs to be taken into account or is an intrinsic part of the analysis. Many, if not all, of the most interesting results from RHIC arise from evaluating how various observables evolve with increasing number of nucleon participants and/or collisions. Experience has also shown that multiple centrality measures, and in particular finding the centrality using a different region of phase space from the observable of interest, are of great value. Centrality determination with the ZDC uses the fact that the number of spectator neutrons is correlated with the impact parameter. Because the CMS ZDC will have a significantly wider acceptance than was the case at RHIC, it is conceivable that very forward π^0 s could contaminate this measurement. However, the expectation based on RHIC results from BRAHMS and PHOBOS is that the yields of both neutrons and π^0 s will decrease with decreasing impact parameter for the most central events. Also, the presence of an electromagnetic section will assist in disentangling the differing contributions to the signal.

The study of elliptic flow, as well as the dependence of other observables on the reaction plane of the initial collision, has contributed enormously to the understanding of RHIC results. As with centrality, multiple measures spread over a range of rapidity are essential. It has been shown at RHIC that the reaction plane derived from the spectator neutrons is in good agreement with other methods [52]. As was the case at RHIC, a small detector that sits near the maximum of the hadronic shower will be used for this measurement. Since this detector will only be used for heavy ions, it does not have to be particularly radiation hard. When in use, this detector will be installed in front of the hadronic calorimeter which will be moved back by 5 cm. For the LHC, a reduction in the strength of the flow signal is likely, perhaps signaling the transition to a weakly interacting quark gluon plasma, so this is both an interesting and challenging measurement.

In addition to their critical role in the determination of centrality and reaction plane, the ZDCs can contribute to other aspects of the physics program for both heavy ion and proton collisions. Ultra-peripheral collisions of highly charged and highly relativistic ions involve an enormous flux of high energy virtual photons, thus offering the opportunity to study several very exotic processes in both QCD and QED. The signals from the ZDC can be used to trigger on such interactions, as well as diffractive events in p+p. In addition, the ZDC and other forward calorimetry in CMS contribute significantly to the hermeticity for determining missing energy. Reaching outside heavy ion and high energy physics, proton-nucleus collisions at LHC energies are equivalent to the interaction of extremely high energy cosmic rays with the Earth's atmosphere. Controlled measurements with known energy and centrality at the LHC will contribute significantly to the quantitative understanding of detectors studying these extreme cosmic rays.

D.2 ZDC Maintenance Costs

The radiation load on the phototubes is comparable to that of the tubes for the HF calorimeter. However, it is possible that a beam accident could damage phototubes of one or the other ZDC. For this reason, funds for ZDC maintenance are included in the operating costs discussed in Appendix B. The activation of the ZDC has been simulated and the residual dose rate on contact at the outer surface of the TAN, which houses the ZDC, is below 0.2 mSv/hr about a day after shutdown. Even after 20 years of continuous operation it should be possible to work at the TAN within a few weeks of shutdown. No significant disposal costs for the ZDC at the end of LHC operations are expected.

D.3 Status of the ZDC Project

The ZDC group has been assigned time in early August, 2006, for a beam test at CERN. A prototype of the calorimeter is being assembled and will be shipped from Kansas in late July, 2006. The actual CMS readout electronics will be used. In consultation with the LHC integration group, many of the details of how to install the ZDC were worked out earlier this year. For example, the cable paths and installation timetable have been set. A dummy ZDC was made from plastic and successfully inserted into the two TANs that will be used for CMS. Finally, a wooden TAN has been constructed to help sort out installation issues on the surface.

A simulation of the ZDC is running under the new CMS software framework. The first round of “validation” compared a simulation of the RHIC ZDCs to existing test beam results for that detector [50]. The second round will be comparison of test beam data to be taken in August to the simulation of the prototype calorimeter. Eventually, data for single 2.7 TeV neutrons from heavy ion beams will be studied. In order to make the simulation fast enough to be useful for flow and centrality studies, a parametrized Monte Carlo, utilizing shower libraries, is being developed.

References

- [1] I. Vitev and M. Gyulassy, Phys. Rev. Lett. **89**, 252301 (2002); I. Vitev, nucl-th/0308028.
- [2] T.D. Lee and M. Gyulassy, nucl-th/0403032.
- [3] J. F. Gunion and R. Vogt, Nucl. Phys. **B492**, 301 (1997).
- [4] M. Gyulassy, P. Levai and I. Vitev, Phys. Rev. Lett. **85**, 5535 (2000).
- [5] U.A. Wiedemann, Nucl. Phys. **B588**, 303 (2000).
- [6] M. Gyulassy and M. Plumer, Phys. Lett. B **243**, 432 (1990).
- [7] M. Plumer, M. Gyulassy and X.-N. Wang, Nucl. Phys. **A590**, 511c (1995).
- [8] C. Adler *et al.* [STAR Collaboration], Phys. Rev. Lett. **90**, 082302 (2003).
- [9] X.-N. Wang, Z. Huang, and I. Sarcevic, Phys. Rev. Lett. **77**, 231 (1996).
- [10] V. Kartvelishvili, R. Kvatadze and R. Shanidze, Phys. Lett. B **356**, 589 (1995).
- [11] R. Vogt, Phys. Rev. C **64**, 044901 (2001).
- [12] G. Baur *et al.*, CMS/2000-60, 2000.
- [13] X.N. Wang, Phys. Rep. **280**, 287 (1997).
- [14] A. Accardi *et al.*, hep-ph/0308248.
- [15] A. Accardi *et al.*, hep-ph/0310274.
- [16] R. Vogt, private communication.
- [17] M. Bedjidian *et al.*, hep-ph/0311048.
- [18] D. Brandt, LHC Project Report 450, 2000.
- [19] K.J. Eskola, V.J. Kolhinen and R. Vogt, Phys. Lett. B **582**, 157 (2004).
- [20] L. McLerran and R. Venugopalan, Phys. Rev. D **49**, 2233 (1994); Phys. Rev. D **59**, 094002 (1999); E. Iancu, A. Leonidov and L. D. McLerran, Nucl. Phys. **A692**, 583 (2001), and refs. therein.
- [21] B.B. Back *et al.*, [PHOBOS Collaboration], Phys. Rev. Lett. **91**, 052303 (2003).
- [22] I. Arsene *et al.* [BRAHMS Collaboration], nucl-ex/0401025.
- [23] I. G. Bearden *et al.*, [BRAHMS Collaboration], Phys. Rev. Lett. **88**, 202301 (2002).
- [24] I. G. Bearden *et al.*, [BRAHMS Collaboration], Phys. Lett. **B607**, 42-50 (2005).
- [25] B.B. Back *et al.*, [PHOBOS Collaboration], Phys. Rev. Lett. **94**, 122303 (2005) and nucl-ex/0511045.
- [26] J. Adams *et al.* [STAR Collaboration], Phys.Rev. **C73**, 034906 (2006) and nucl-ex/0511026.
- [27] M. Murray, [BRAHMS Collaboration], nucl-ex/0602015.
- [28] L. A. Stiles and M. Murray, nucl-ex/0601039.
- [29] I. G. Bearden *et al.* [BRAHMS Collaboration], nucl-ex/0312023.
- [30] I. G. Bearden *et al.* [BRAHMS Collaboration], nucl-ex/0403050.
- [31] M. Murray *et al.* [BRAHMS Collaboration], nucl-ex/0404007.
- [32] M. Strikman, R. Vogt, and S. White, Phys. Rev. Lett. **96**, 082001 (2006).

- [33] I. Arsene *et al.* [BRAHMS Collaboration], nucl-ex/0403005.
- [34] M. Chiu, A. Denisov, E. Garcia, J. Katzy, A. Makeev, M. Murray and S. White, Phys. Rev. Lett. **89**, 012302 (2002) and nucl-ex/0109018.
- [35] D. Alton *et al.*, Phys. Rev. D**56**, 5301 (1997).
- [36] M. Greiner, M. Vidovic, Ch. Hofmann, A. Schäfer and G. Soff, Phys. Rev. C **51**, 911 (1995); N. Baron and G. Baur, Phys. Rev. C **48**, 1999 (1993); S.R. Klein, J. Nystrand and R. Vogt, Phys. Rev. C **66**, 044906 (2002).
- [37] S.R. Klein, J. Nystrand and R. Vogt, Eur. Phys. J. C **21**, 563 (2001).
- [38] F. Gelis and A. Peshier, Nucl. Phys. A**697**, 879 (2002).
- [39] CMS HCAL Design Report, CERN/LHCC 97-31.
- [40] CMS MUON Design Report, CERN/LHCC 97-32.
- [41] CMS ECAL Design Report, CERN/LHCC 97-33.
- [42] CMS Tracker Design Report, CERN/LHCC 98-6.
- [43] A. Angelis and A.D. Panagiotou, Journ. Phys. G: Nuclear & Particle Physics **23**, 18 (1997).
- [44] G. Matthiae, TOTEM Technical Proposal, CERN/LHC 99-7 (1999).
- [45] The four RHIC collaborations, BRHAMS, PHENIX, PHOBOS and STAR have published papers that summarize the first years of RHIC physics in Nucl. Phys. A**757** (2005).
- [46] B.B. Back *et al.*, [PHOBOS Collaboration], Nucl. Phys. A**757**, 28 (2005).
- [47] I.P. Lokhtin, A.M. Snigirev, Eur. Phys. J. C**46**, 211 (2006).
- [48] T. Sjostrand, P. Eden, C. Friberg, L. Lonnblad, G. Miu, S. Mrenna and E. Norrbin, Computer Phys. Commun. **135**, 238 (2001); L. Lonnblad, S. Mrenna, T. Sjostrand and P. Skands, LU TP 03-38, hep-ph/0308153.
- [49] M. Bedjidian and O. Kodolova, CMS Note “Quarkonia measurements in heavy ion collisions in CMS”.
- [50] C. Adler, A. Denisov, E. Garcia, M. Murray, H. Strobele, S. White, Nucl. Instr. Meth. A**470** (2001) 488.
- [51] T. Zimmerman and J.R. Hoff, IEEE Journal of Solid State Circuits **39**, 895 (2004).
- [52] J. Adams *et al.*, [STAR Collaboration], Phys. Rev. C**73**, 034903 (2006).
- [53] CMS The Computing Project Technical Design Report, CERN-LHCC-2005-023.